BVQ Documentation

BVQ user's guide

Author: BVQ Team

Date: 11/29/2013

TABLE OF CONTENTS

1 User's guide - BVQ GUI 5

1.1 How to start the GUI 6
1.1.1 Cluster selection 6

1.2 Main menu 7
1.2.1 Menu "File" 7
1.2.2 Menu "object lists" 9
1.2.3 Menu "favorites" 10
1.2.4 Menu "?" 10

1.3 Favorites and tabs 19
1.3.1 The control functions 19
1.3.2 Favorites 21

1.4 Scanner status bar 23

1.5 Topology view 23
1.5.1 Treemap 25
1.5.2 Topology - Object path - Topology navigator 35
1.5.3 Topology - aspects 40
1.5.4 Topology - color schemes 47
1.5.5 Topology - search function 68

1.6 Performance view 70
1.6.1 Main panels 71
1.6.2 Walkthrough 72
1.6.3 Call of performance view 72
1.6.4 Selection of the time period 74
1.6.5 Loading and drawing of the graph 75
1.6.6 Scaling of the value axes 75
1.6.7 Options dialog for objects and metrics 76
1.6.8 Chart area 80
1.6.9 Performance metrics 82

1.7 Tables view 104
1.7.1 Filtering 105
1.7.2 Sorting the table 106
1.7.3 Search 107
1.7.4 Other functions 108

1.8 Object detail view 108
1.8.1 Calling the object detail view 109
1.8.2 Overview object detail view 109

1.8.3 Categories 109
1.8.4 Information groups 115

1.9 Key figures of the BVQ storage environment 115
1.9.1 Storage capacities 116
1.9.2 Storage costs and prices 124
1.9.3 Object counter 124
1.9.4 Performance indicators 125

1.10 Object types 128
1.10.1 Overview object types 129
1.10.2 BVQ sites 130
1.10.3 Storage providers 130
1.10.4 SVC virtualization 131
1.10.5 SVC advanced copy services 131
1.10.6 Storage consumers 131
1.10.7 SAN connections and ports 132
1.10.8 BVQ organizational units 133
1.10.9 BVQ placement rules 133
1.10.10 BVQ service level management 133
1.10.11 BVQ - VMware virtualization 133
1.10.12 BVQ organizational units 134
1.10.13 BVQ placement rules 137
1.10.14 BVQ service level management 142
1.10.15 BVQ sites 145
1.10.16 BVQ - VMware virtualization 146
1.10.17 SAN connections and ports 147
1.10.18 Storage consumers 148
1.10.19 Storage providers 149
1.10.20 SVC advanced copy services 154
1.10.21 SVC virtualization 154

1.11 Object editors 155
1.11.1 Object-editors 158

1.12 Use cases 160
1.12.1 Qualification of the performance edition and the single extension-packages 160
1.12.2 Usage of the treemap for the optimization of storage costs 212
1.12.3 Capacity efficiency 214

2 User's guide - BVQ reporting 221

2.1 Reports overview 221
2.1.1 Standard reports 221
2.1.2 Premium reports 221
2.1.3 Customer-orientated reports 221

2.2 Creating reports 222

2.3 ODBC access to the BVQ reporting scheme 223
2.3.1 Content 223
2.3.2 Overview 224
2.3.3 Download 224
2.3.4 Installation 224
2.3.5 Configuration 225
2.3.6 Procedure on 64 bit systems 230
2.3.7 Usage in MS Excel 2007 230
2.3.8 More information 233

2.4 BVQ database view scheme for report - scheme level 1.0.0.0 234

3 Annex 245

3.1 Most important links 245

BVQ_Version	3.0

1 USER'S GUIDE - BVQ GUI

The BVQ GUI is used for visualization and configuration of objects in the maintained storage environment.

Abbildung 1 GUI controls

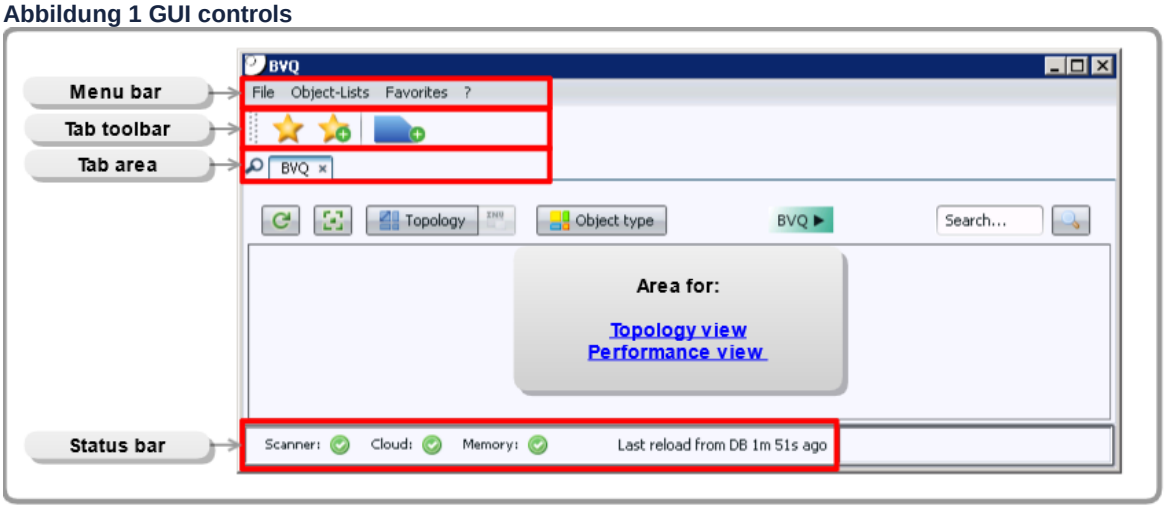

The GUI provides you with an overview of the topology, performance, costs and compliance of your storage infrastructure.

After launching, the GUI connects with the BVQ database (default: IP port 50000) and loads the most recent state of the selected SVC cluster(s).

The GUI consists of the following main modules:

- Common controls
  - Main menu
  - Favorites and tabs
  - Scanner status bar
- Views
  - Topology view - Treemaps
  - Performance view - System load charts
  - Tables view - Objects in a tabular format
  - Object detail view - Attributes of an individual object in a tabular format
- Object editors - Configuration of object attributes and relations

1.1 How to start the GUI

BVQ can be started in two different ways: via the Windows start menu or the desktop icon.

1. Start via the Windows start menu**: start → programs → SVA → BVQ → BVQ GUI → launch BVQ GUI**

Abbildung 2 GUI - Starting via Windows start menu

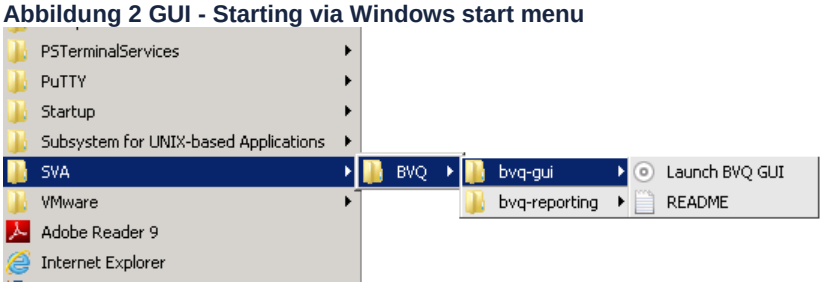

2. Or double-click the link on your desktop: **launch BVQ GUI**.

Abbildung 3 GUI - Starting via desktop icon

1.1.1 Cluster selection

If more than one system (SVC cluster / Storwize) is administered with BVQ, the cluster choice dialogue appears after the start of the BVQ GUI.
The charging time and the memory consumption of the GUI can be reduced by selecting less clusters.

Abbildung 4 Cluster selection dialogue

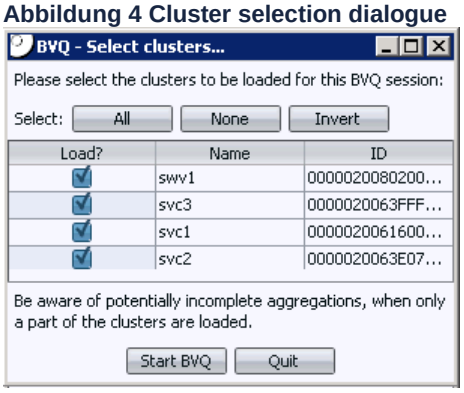

1.2 Main menu

In the main menu of the BVQ GUI functions are found, which are working across different views.

Tabelle 1 Main menu - function groups

File	General functions
Object-Lists	Call of the table views
Favorites	Call and add favorites
?	Miscellaneous information about the BVQ system (version, licence, storage usage)

1.2.1 Menu "File"

Main Menu → File provides access to various functions.

Abbildung 5 Main menu - File

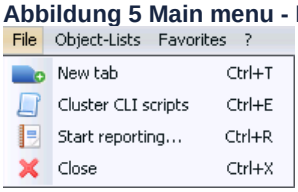

Tabelle 2 Main menu file: functions

New tab	Opens a new tab with a blank topology view
Cluster CLI scripts	Opens the window *Cluster CLI scripts*
Start reporting ...	Starts the BVQ-Reporting GUI
Close	Closes the GUI

1.2.1.1 Cluster CLI Scripts

"Cluster CLI scripts" collects SVC commands which were generated by drag & drop.

Abbildung 6 Cluster CLI scripts dialog

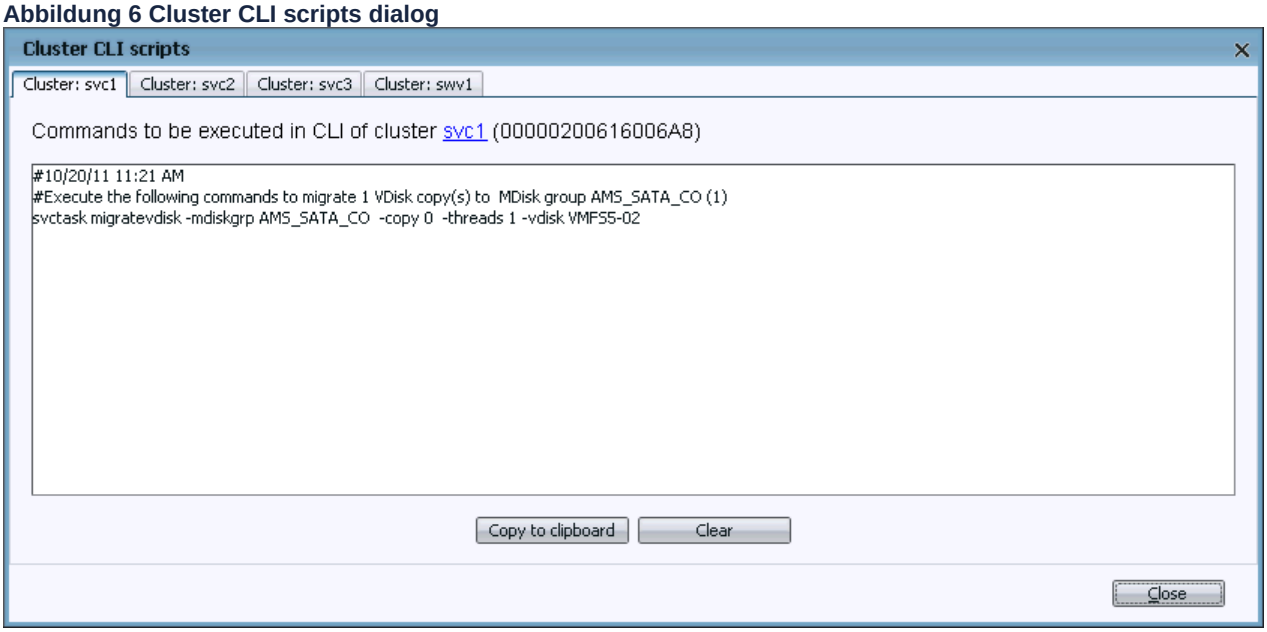
Cluster CLI scripts
Cluster: svc1
Cluster: svc2
Cluster: svc3
Cluster: swv1
Commands to be executed in CLI of cluster svc1 (00000200616006A8)
#10/20/11 11:21 AM
#Execute the following commands to migrate 1 VDisk copy(s) to MDisk group AMS_SATA_CO (1)
svctask migratevdisk -mdiskgrp AMS_SATA_CO -copy 0 -threads 1 -vdisk VMFS5-02
Copy to clipboard
Clear
Close

The commands can be subsequently sighted, edited and performed in the SVC CLI.
The script is not lost even after the dialog is closed.

Tabelle 3 Cluster CLI scripts - functions of the dialog

Cluster tabs	A separate dialog is conducted for each cluster.
Hyperlink to SVC CLI	Calls a SSH session for this cluster.
Text field	Contains by drag & drop generated SVC commands including a time value. Here the commands can be edited later.
Copy to clipboard	Copies the commands from the text field into the system clipboard to run them, for example, in a SSH session.
Clear	Deletes the content of the text field.
Close	Closes cluster CLI scripts

Call ing the SVC CLI by using SSH client " PuTTY "

The BVQ GUI calls the SVC CLI by using the SSH client "PuTTY" with the following command:
`putty -load [clustername]`

For the call´s correct functioning the following conditions must be met:

- PuTTY has to be installed (download available at http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html)
- The command `putty` has to be included in the system search path `%PATH%` (typical for Microsoft Windows Server 2008 x64: `C:\PROGRA~2\PuTTY`).
- In PuTTY appropriate sessions have to be configured with the same cluster name.

If the command cannot be called, a message appears referencing the used system command.
e.g.: "putty -load [clustername] not successfull"

1.2.2 Menu "object lists"

Mein menu → object-lists provides access to the tables view for all editable objects.

The object types in this menu are classified in the same categories as in the path navigator.

Abbildung 7 Main menu - object lists

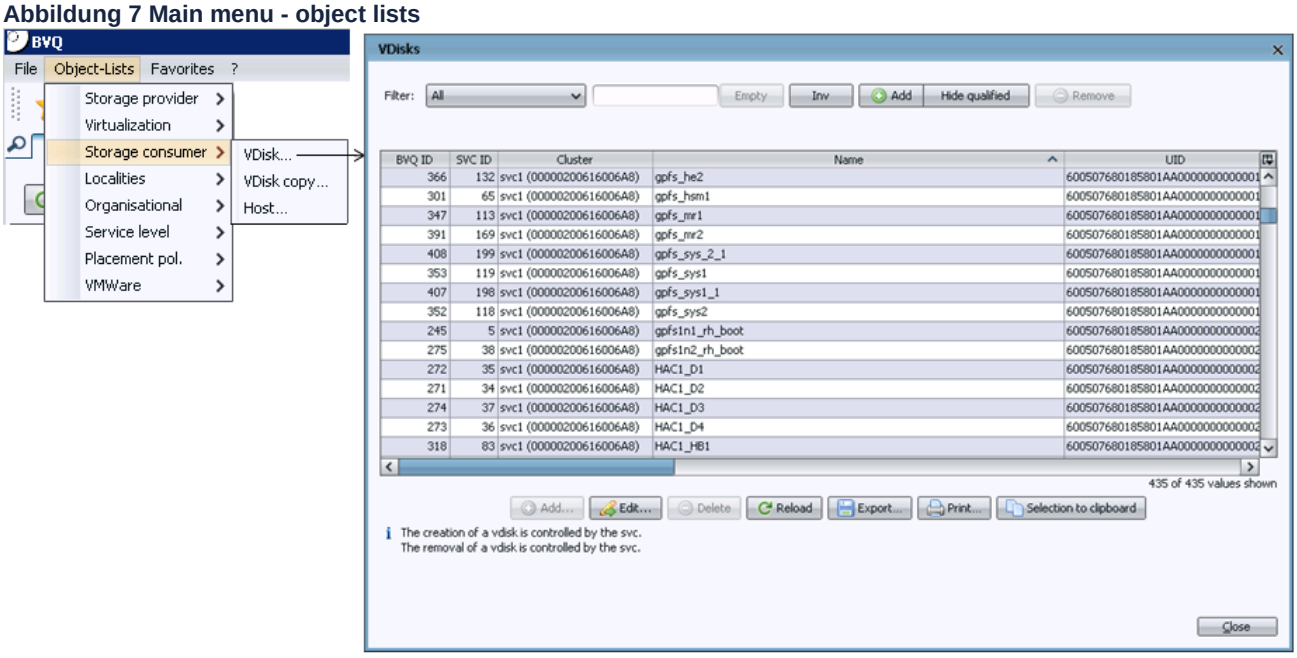

⚠ Depending on the internal BVQ license key some object types and their menu items are disabled.

1.2.3 Menu "favorites"

Main menu → favorites opens the favorites selection which allows access to stored topology views and performance views.

Abbildung 8 Main menu - favorites

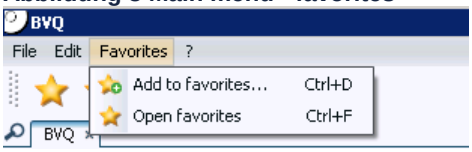

Tabelle 4 Main menu - favorites functions

Main menu functions – favorites	
Add to favorites ...	Allows you to save all views, performance graphs included.
Open favorites	Opens selection of favorites.

1.2.4 Menu "?"

Main menu →? provides access to functions for license managing, information about the GUI memory consumption, scanner status and code version.

Abbildung 9 Main menu - ?

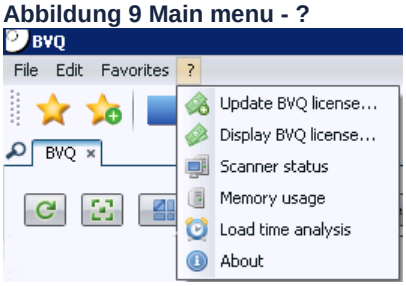

Tabelle 5 Main menu - ? - functions

Update BVQ license ...	Allows the import of a BVQ license key file .
Display BVQ license ...	Displays the current licensing and its compliance in a status window.

Scanner status	The scanner status-window shows the current state of BVQ scan processes.
Memory usage	The memory usage window shows the current memory usage of the GUI.
Load time analysis	The window *load time analysis* shows the load times of each BVQ objects when starting the GUI.
About	The *about* window displays product-related data, such as version number, contact information, legal information etc.

1.2.4.1 Management of license keys

Description of BVQ license keys´ management (BVQ license key files).

The license key management functions are accessed via the menu "?" in the main menu:

Abbildung 10 Main menu - ? managing license keys

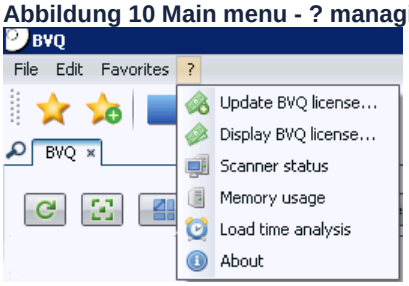

Tabelle 6 License key - management functions

	Update BVQ license...	Import / update new license key file
	Display BVQ license...	View current license key

Overview of licensed feature packs in the license table

Abbildung 11 License table: status window

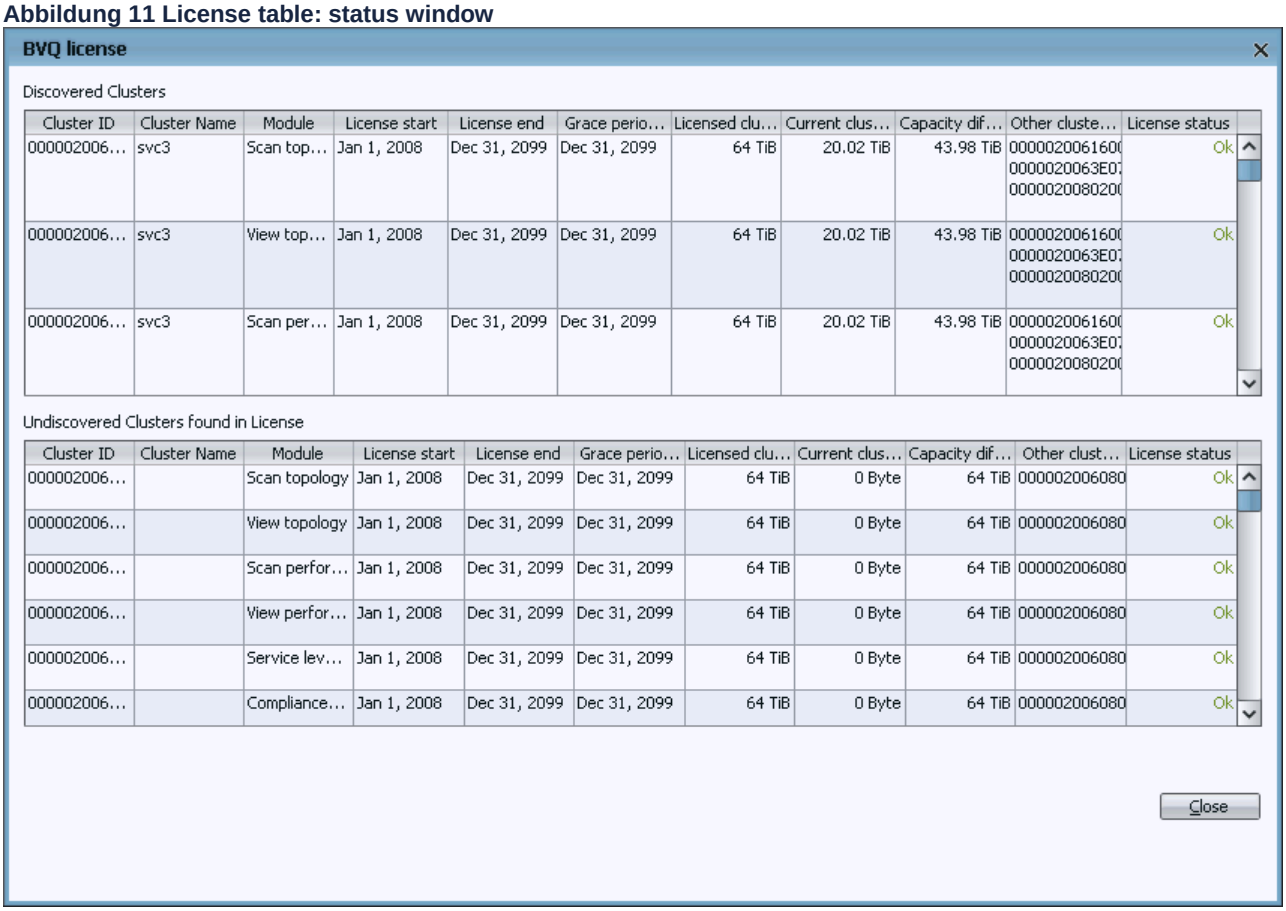
BVQ license
Discovered Clusters
Cluster ID | Cluster Name | Module | License start | License end | Grace perio... | Licensed clu... | Current clus... | Capacity dif... | Other cluste... | License status
000002006... | svc3 | Scan top... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 20.02 TiB | 43.98 TiB | 0000020061600 0000020063E07 0000020080200 | Ok
000002006... | svc3 | View top... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 20.02 TiB | 43.98 TiB | 0000020061600 0000020063E07 0000020080200 | Ok
000002006... | svc3 | Scan per... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 20.02 TiB | 43.98 TiB | 0000020061600 0000020063E07 0000020080200 | Ok
Undiscovered Clusters found in License
Cluster ID | Cluster Name | Module | License start | License end | Grace perio... | Licensed clu... | Current clus... | Capacity dif... | Other clust... | License status
000002006... | | Scan topology | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 0 Byte | 64 TiB | 000002006080 | Ok
000002006... | | View topology | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 0 Byte | 64 TiB | 000002006080 | Ok
000002006... | | Scan perfor... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 0 Byte | 64 TiB | 000002006080 | Ok
000002006... | | View perfor... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 0 Byte | 64 TiB | 000002006080 | Ok
000002006... | | Service lev... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 0 Byte | 64 TiB | 000002006080 | Ok
000002006... | | Compliance... | Jan 1, 2008 | Dec 31, 2099 | Dec 31, 2099 | 64 TiB | 0 Byte | 64 TiB | 000002006080 | Ok
Close

Tabelle 7 License table: chapter discovered / undiscovered cluster

Tables	
Discovered clusters	Table of licensed and already recognized SVC/V7000 clusters
Undiscovered clusters found in license	Table of licensed but not yet recognized SVC/V7000 clusters

Tabelle 8 License table: explanation of columns

Columns	
Cluster ID	Corresponds to the SVC/V7000 cluster **alias** ID (s. `id_alias` from `svcinfo lscluster`)
Cluster Name	Name of the cluster

Module	Name of the function module
License start	License valid from
License end	License valid until
Grace period end	License renewal period ending on
Licensed cluster pool capacity	Total licensed capacity of all clusters in the pool (in TiB)
Current cluster pool capacity	Currently used capacity (in TiB)
Capacity difference	Unused licensed capacity
Other clusters in pool	List of other licensed BVQ clusters
License status	Status of the module

Tabelle 9 License table: license statuses for each module

Status	Description	BVQ reaction
OK	Active and in use license compliant	Normal operation
Grace period	Within the extension period	Status window is displayed
Expired	Violation: license and renewal period expired	Module inactive
Size exceeded	Violation: licensed capacity exceeded	Status window is displayed
Missing	License file contains no entry for this module	Module inactive
Invalid	Faulty entry in the license file for this module	Module inactive
Not licensed	Not active and not licensed	Module inactive

⚠ Red status messages urgently require a user reaction.

In case of license violations the license table is automatically opened in the foreground:

- Shortly after the start of the GUI
- At regular intervals (approximately every 12 hours)

An overview about the licensable feature packs are available in the annex and BVQ license matrix.

1.2.4.2 Scanner status

The scanner status-window shows the current state of BVQ scan processes.

Abbildung 12 Scanner status-window

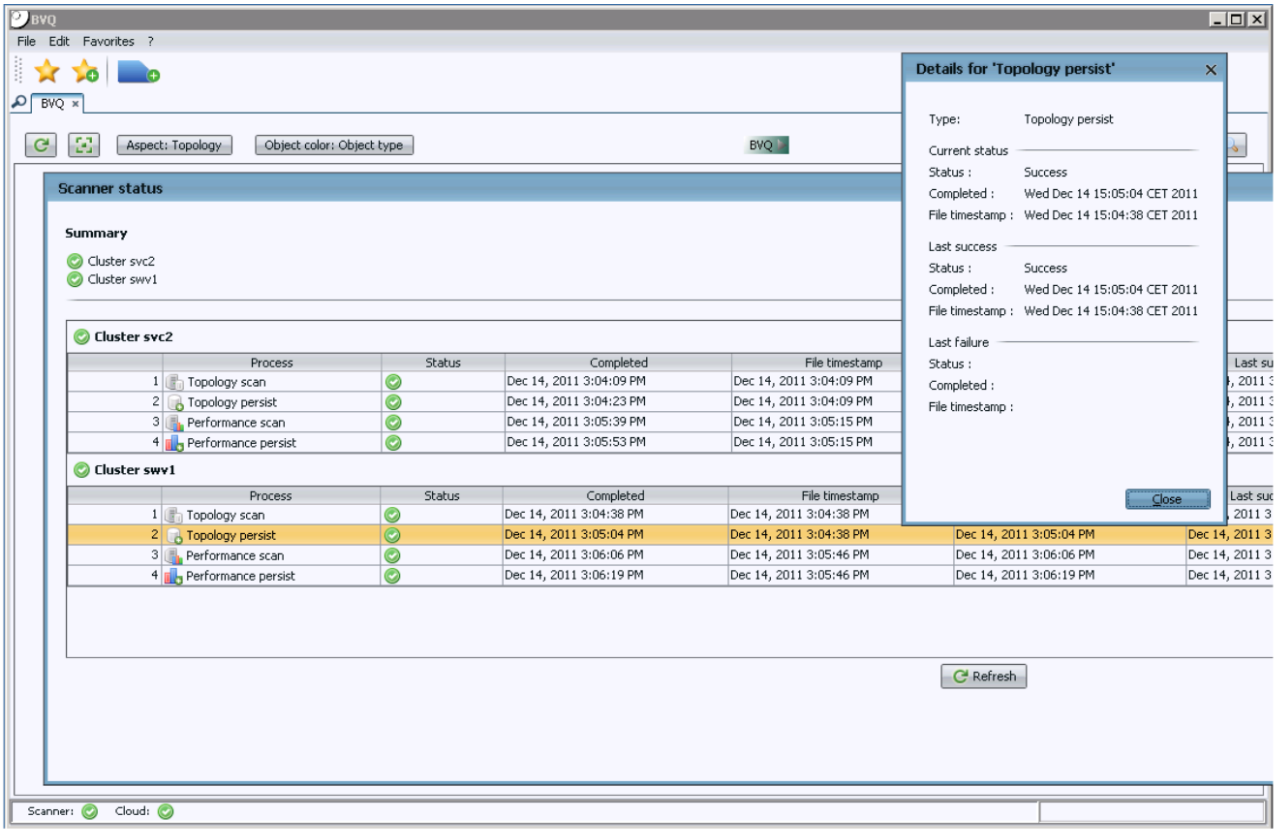
BVQ
File Edit Favorites ?
BVQ
Aspect: Topology
Object color: Object type
BVQ
Scanner status
Summary
Cluster svc2
Cluster swv1
Cluster svc2
Process Status Completed File timestamp
1 Topology scan Dec 14, 2011 3:04:09 PM Dec 14, 2011 3:04:09 PM
2 Topology persist Dec 14, 2011 3:04:23 PM Dec 14, 2011 3:04:09 PM
3 Performance scan Dec 14, 2011 3:05:39 PM Dec 14, 2011 3:05:15 PM
4 Performance persist Dec 14, 2011 3:05:53 PM Dec 14, 2011 3:05:15 PM
Cluster swv1
Process Status Completed File timestamp
1 Topology scan Dec 14, 2011 3:04:38 PM Dec 14, 2011 3:04:38 PM
2 Topology persist Dec 14, 2011 3:05:04 PM Dec 14, 2011 3:04:38 PM Dec 14, 2011 3:05:04 PM Dec 14, 2011 3
3 Performance scan Dec 14, 2011 3:06:06 PM Dec 14, 2011 3:05:46 PM Dec 14, 2011 3:06:06 PM Dec 14, 2011 3
4 Performance persist Dec 14, 2011 3:06:19 PM Dec 14, 2011 3:05:46 PM Dec 14, 2011 3:06:19 PM Dec 14, 2011 3
Refresh
Details for 'Topology persist'
Type: Topology persist
Current status
Status : Success
Completed : Wed Dec 14 15:05:04 CET 2011
File timestamp : Wed Dec 14 15:04:38 CET 2011
Last success
Status : Success
Completed : Wed Dec 14 15:05:04 CET 2011
File timestamp : Wed Dec 14 15:04:38 CET 2011
Last failure
Status :
Completed :
File timestamp :
Close
Scanner:
Cloud:

Each of BVQ managed storage system (SVC, V7000) is observed by a scanner.
BVQ divides the scan process into four sub-processes:

Tabelle 10 Scan sub-processes

1. Topology scan	Reads configuration information from the system memory and writes it to an XML file into a local queue
2. Topology persist	Reads configuration information from the queue and writes them to the database
3. Performance scan	Copies performance statistic files from the storage system and writes them to a local queue
4. Performance	Reads performance statistics files from the queue and writes them to the

persist	database

Tabelle 11 Scanner status - columns of the tables

Process	Name of the sub-process
Status	Displays the states OK , Warning, Error
Completed	Displays when the last cycle was completed
File timestamp	Displays the timestamp of the last processed topology or performance statistics file
Last successfull completion	Displays the timestamp of the last successful cycle
Last successfull file timestamp	Displays the timestamp of the last successfully processed topology or performance statistics file
Last failed completion	Displays the timestamp of the last failed cycle
Last failed file timestamp	Displays the timestamp of the last failed topology or performance statistics file

Function refresh

By clicking the refresh button a re-query of the database is initiated, last scan and persist cycles can be recorded and evaluated in the scanner status dialog.

Summary of conditions

A summary of the conditions can be found at the top of the status window and in the scanner status bar. Always the most severe status is reported in the summary.

1.2.4.3 Memory usage

The memory usage window shows the current memory usage of the GUI.

Abbildung 13 Window memory usage

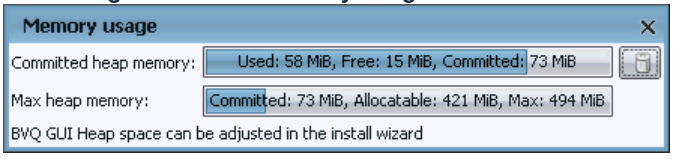
Memory usage
Committed heap memory:
Used: 58 MiB, Free: 15 MiB, Committed: 73 MiB
Max heap memory:
Committed: 73 MiB, Allocatable: 421 MiB, Max: 494 MiB
BVQ GUI Heap space can be adjusted in the install wizard

Tabelle 12 Inofrmation about GUI memory usage

Used	Amount of allocated memory
Free	Amount of free memory in relation to committed
Committed	Size of the reserved memory for the GUI
Allocatable	Size of the remaining available memory in relation to Max
Max	Size of the maximum available memory for the GUI process*
	Switch to trigger a Java runtime garbage collection

* The maximum amount of memory depends on the value of "Max Heap Size", which can be configured in the GUI installation.

1.2.4.4 Load time analysis

The window *load time analysis* shows the load times of each BVQ objects when starting the GUI.

Abbildung 14 Window load time analysis

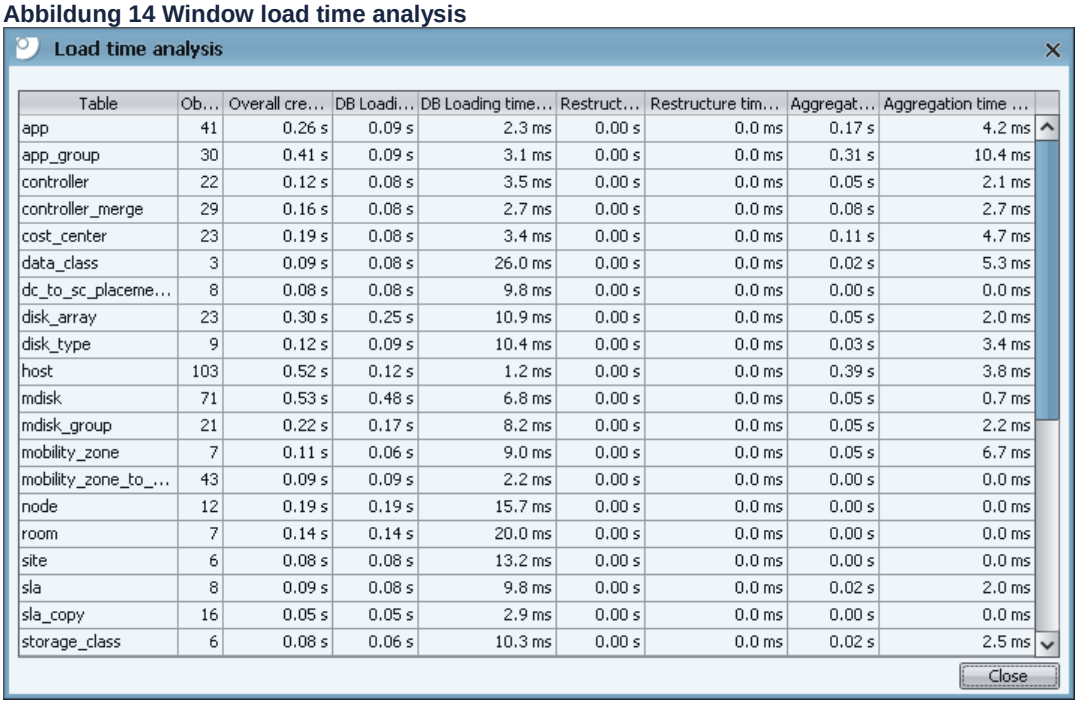

Table	Ob...	Overall cre...	DB Loadi...	DB Loading time...	Restruct...	Restructure tim...	Aggregat...	Aggregation time ...
app	41	0.26 s	0.09 s	2.3 ms	0.00 s	0.0 ms	0.17 s	4.2 ms
app_group	30	0.41 s	0.09 s	3.1 ms	0.00 s	0.0 ms	0.31 s	10.4 ms
controller	22	0.12 s	0.08 s	3.5 ms	0.00 s	0.0 ms	0.05 s	2.1 ms
controller_merge	29	0.16 s	0.08 s	2.7 ms	0.00 s	0.0 ms	0.08 s	2.7 ms
cost_center	23	0.19 s	0.08 s	3.4 ms	0.00 s	0.0 ms	0.11 s	4.7 ms
data_class	3	0.09 s	0.08 s	26.0 ms	0.00 s	0.0 ms	0.02 s	5.3 ms
dc_to_sc_placeme...	8	0.08 s	0.08 s	9.8 ms	0.00 s	0.0 ms	0.00 s	0.0 ms
disk_array	23	0.30 s	0.25 s	10.9 ms	0.00 s	0.0 ms	0.05 s	2.0 ms
disk_type	9	0.12 s	0.09 s	10.4 ms	0.00 s	0.0 ms	0.03 s	3.4 ms
host	103	0.52 s	0.12 s	1.2 ms	0.00 s	0.0 ms	0.39 s	3.8 ms
mdisk	71	0.53 s	0.48 s	6.8 ms	0.00 s	0.0 ms	0.05 s	0.7 ms
mdisk_group	21	0.22 s	0.17 s	8.2 ms	0.00 s	0.0 ms	0.05 s	2.2 ms
mobility_zone	7	0.11 s	0.06 s	9.0 ms	0.00 s	0.0 ms	0.05 s	6.7 ms
mobility_zone_to_...	43	0.09 s	0.09 s	2.2 ms	0.00 s	0.0 ms	0.00 s	0.0 ms
node	12	0.19 s	0.19 s	15.7 ms	0.00 s	0.0 ms	0.00 s	0.0 ms
room	7	0.14 s	0.14 s	20.0 ms	0.00 s	0.0 ms	0.00 s	0.0 ms
site	6	0.08 s	0.08 s	13.2 ms	0.00 s	0.0 ms	0.00 s	0.0 ms
sla	8	0.09 s	0.08 s	9.8 ms	0.00 s	0.0 ms	0.02 s	2.0 ms
sla_copy	16	0.05 s	0.05 s	2.9 ms	0.00 s	0.0 ms	0.00 s	0.0 ms
storage_class	6	0.08 s	0.06 s	10.3 ms	0.00 s	0.0 ms	0.02 s	2.5 ms

Tabelle 13 Load time analysis - Meaning of the table columns

Table	Displays the corresponding BVQ table, from which the BVQ object can be

	derived
Objects	Displays the number of loaded objects from the respective table
Overall creation time	Displays the total duration for the creating process of the objects
DB loading time	Displays the total time of all objects of the appropriate type from the DB
DB loading time per object	Displays the average duration of loading a single object of the respective type
Restructure time	Displays the total arrangement´s duration for all objects of the appropriate type into the BVQ cloud structure
Restructure time per object	Displays the average arrangement´s duration of a single object into the BVQ cloud structure
Aggregation time	Displays the time of the aggregation calculations in the cloud of BVQ
Aggregation time per object	Displays the time of the aggregation calculations per object in the cloud of BVQ

1.2.4.5 About dialog

The *about* window displays product-related data, such as version number, contact information, legal information etc.

Abbildung 15 About window

About
BVQ
Business Volume Qualicision
SVA
Serial number: BVQ00010101
Version: 2.6.1
Build number: 19
DB schema version: 2.6.0.3
JRE version: Java HotSpot(TM) Client VM 1.7.0_05 (Oracle Corporation)
OS: Windows XP Service Pack 3 (32 Bit)
© Copyright: 2006 to 2012 SVA GmbH Wiesbaden Germany - All rights reserved
® Trademarks: BVQ® and BVQ Business Volume Qualicision® are trademarks of SVA GmbH
Patents: All patents reserved
Copy version information to clipboard
Close

Tabelle 14 Information of the about dialogs

Information	
Serial number	Customer´s BVQ serial number
Version	Installed version of the BVQ GUI
Build number	Internal build number of BVQ GUI (for support purposes)
DB scheme version	Installed BVQ DB scheme version of the database
JRE version	Installied Java Runtime Environment version (BVQ 2.6 -> Java 1.7 is required)
OS	Operating system version on which the BVQ GUI is installed

Tabelle 15 Window about - dialog´s functions

Copy version information to clipboard	BVQ GUI version information is copied to the system´s clipboard
Close	Closes the about dialog

1.3 Favorites and tabs

Some views of the GUI are displayed in so-called tabs. Favorites save the settings of these views and facilitate the repeated call.

Tabs and favorites are tied closely.

The view in the currently opened tab can be saved as a favorite, or it can get a new content by selecting a favorite.
More tabs can be created without content, or be opened directly with the saved view of a favorite.

Abbildung 16 Management elements for favorites and tabs

1.3.1 The control functions

The following GUI features are available for the control of tabs and favorites

	Open favorites	Opens the favorites selection
	Add to favorites	Saves the view in the current tab as a user favorite
	Add tab	Creates a new tab with an empty topology view

These functions can be reached with:

Abbildung 17 Tab toolbar

Abbildung 18 Main menu - Favorites

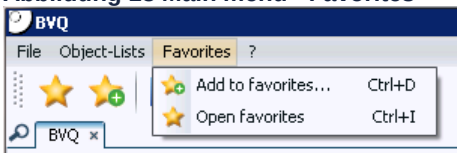

1.3.1.1 Tab overviews

To get a better orientation in a variety of open tabs, two types of tab overviews can be used:

1. Tab overall view - displays all open tabs as thumbnails in a window
2. Tab single preview - shows the selected tab in a thumbnail

Each tab has a title, which is completly displayed in the Tooltip.

Tab total list

Abbildung 19 Calling-up the tab overview

With a click on the little magnifying glass left next to the tabs, the tab overview is opened.
If the mouse pointer is moved over one of the thumbnails, it is zoomed in and can be activated by a mouse click.

Abbildung 20 Example of a tab overview

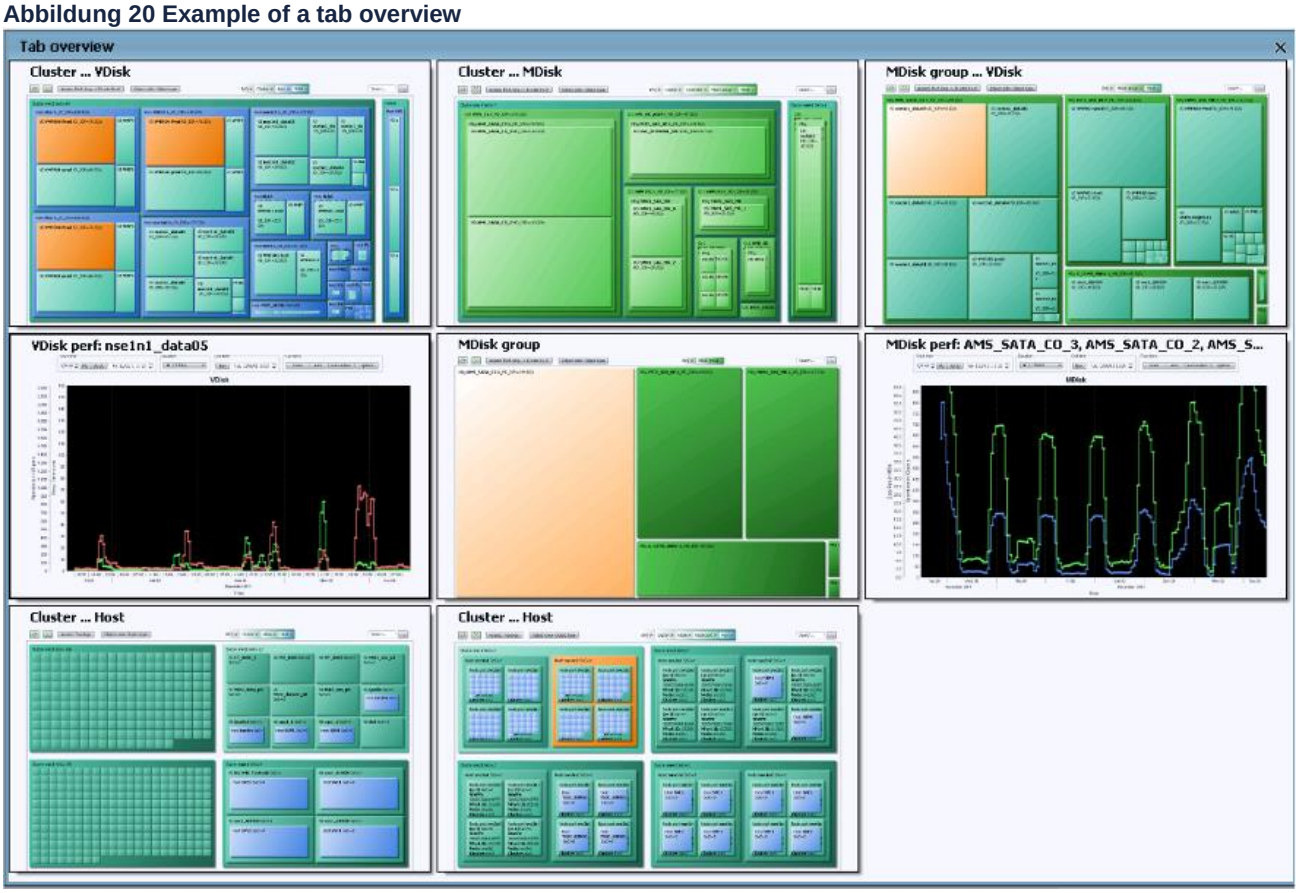

Single preview of tabs

If you rest the mouse pointer slightly longer over a tab, the Tooltip for tabs, which fully displays the title, opens first.
After extra dwell time, a preview image for this tab opens

1.3.2 Favorites

Favorites are saved views of the GUI. They contain views (aspects) from different angles of vision. Either topology or performance views. Favorites are divided due to overview and usage. In BVQ, there are the following divisions:

Predefined favorites	Predefined favorites, delivered with BVQ *
User favorites	Favorites which were saved by a user

* Dependig on the BVQ license, some of these favorites are not available.

1.3.2.1 Favorites selection window

Abbildung 21 Favorites: Overview of stored favorite tabs

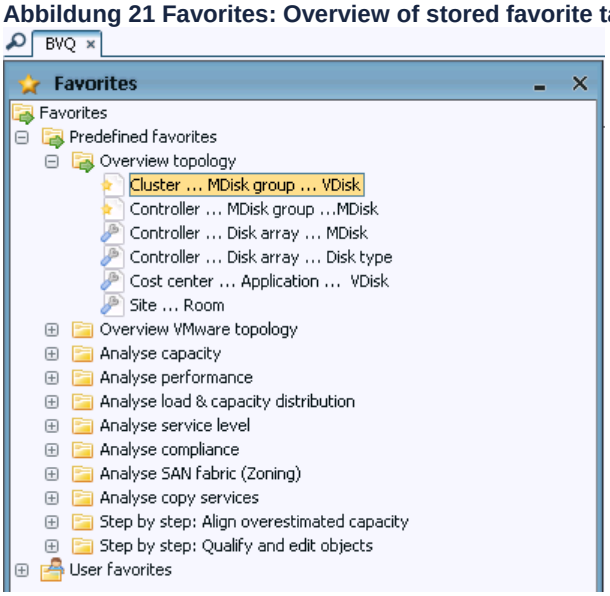

A left click on a favorite opens them in the current tab.

The following functions can be obtained via the context menu of the folder and favorites:

Tabelle 16 Favorites: Functions of the context menu

Rename folder / favorite	Rename of the folder ot favorite
Create new folder	Create a new folder
Delete folder / favorite	Deletion of the folder of favorite
Move folder / favorite to	Relocation into another folder
Open in a new tab	Opens a favorite in a new tab

The favorites are categorized with the following icons:

Tabelle 17 Favorites: Significance of icons

	Predefined topology favorite, which requires more BVQ qualification to deliver useful prompts
	Predefined topology favorite, which needs no further qualification
	Custom topology favorite
	Custom performance favorite

1.4 Scanner status bar

The GUI status bar provides information about the different operating conditions and background activities.

Abbildung 22 Gui status bar

Tabelle 18 Status bar elements

Modules	Purpose			
Scanner	Indicates whether all scanners have provided current data. More information can be found in paragraph scanner status.	All OK	Some delayed	Some corrupt
Cloud	Displays the state of the GUI internal data (networking and aggregations).	Completely calculated	Are calculated	–
Memory	Shows if the GUI can work with sufficient memory.	Sufficient	80% used	100% used
Last reload from DB	Indicates if the data of the BVQ cloud was loaded the last time from the database.			
Progress bar	When background operations of the GUI operations are running, such as calculations or loading operations, a progress bar is displayed righthand on the status bar. By clicking the red button righthand of the bar, the current operations can be canceled manually.			

1.5 Topology view

The topology view displays selected objects of the managed BVQ environment in a treemap

Abbildung 23 Topology view with treemap

Tabelle 19 Topology view: controls

Reload	Loads current data for the topology view and performance indicators from the BVQ DB
Reset zoom	Resets the treemap magnification
Aspekt-Auswahl	The aspects of the topology view allow to refocus the display according to certain priorities.
Farb-Auswahl	Via the color schemes the objects´ areas in the treemap are colored, based on selectable criteria.
Topology Navigator	With the topology navigator several BVQ object types are connected to one path, which determines their hierarchical representation within the treemap.
Topology - search function	Via the search function, objects can be found and specifically selected in the treemap.
Treemap	The treemap is a resource-oriented hierarchical representation of selected objects as a rectangular area.

1.5.1 Treemap

The treemap is a resource-oriented hierarchical representation of selected objects as a rectangular area.

Due to the large number of memory objects in a Storage-Area-Network (SAN,) it is very hard to get an overview of the capacities and needs of the individual object. However, a good overview is desired to decide which store objects are suitbale for an optimization.

A key feature of BVQ is the treemap. This visualization system is aligned on a resource-oriented representation. Its strength is to allow the user a quick and easy acquisition of proportions and ordinal structures at a high number of objects. Treemaps were invented in the early 1990s by Ben Shneiderman.

Each managed object is represented as a rectangular area.

- **Proportional size:** The area size is proportional to an object-specific factor (typical in BVQ: capacity, performance, cost,...).
  The slection of the size factor is made with the help of the *aspect-selector*.
- **Proportional coloring:** The area coloration can represent other object-specific attributes (object type class, conditions for qualification or compliance,...)
  The selection of the dye attribute is done with the help of the *color-selector*.
- **Groupings:** The areas can be interlaced in multiple levels to represent hierarchical groupings. Several small areas, surrounded by a larger area, represent the grouping of the small ones in the large one.
  The selection what type of object is represented in which level is made with the help of the *topology navigator*
- **Tin pie chart:** The areas within a group are sorted by proportional representation according to size: top left the large, right down the smallest.
  In doing so they fill the entire available space in the surrounding area and are thereby similar to the famous pie charts:
  The size of the treemap here has the same function as the area of a circle cut-out in the pie chart (analogy: treemap-> tin pie chart)

The treemap controls are explained in the chapter topology view.

1.5.1.1 Example of a treemap:

Abbildung 24 Treemap-view

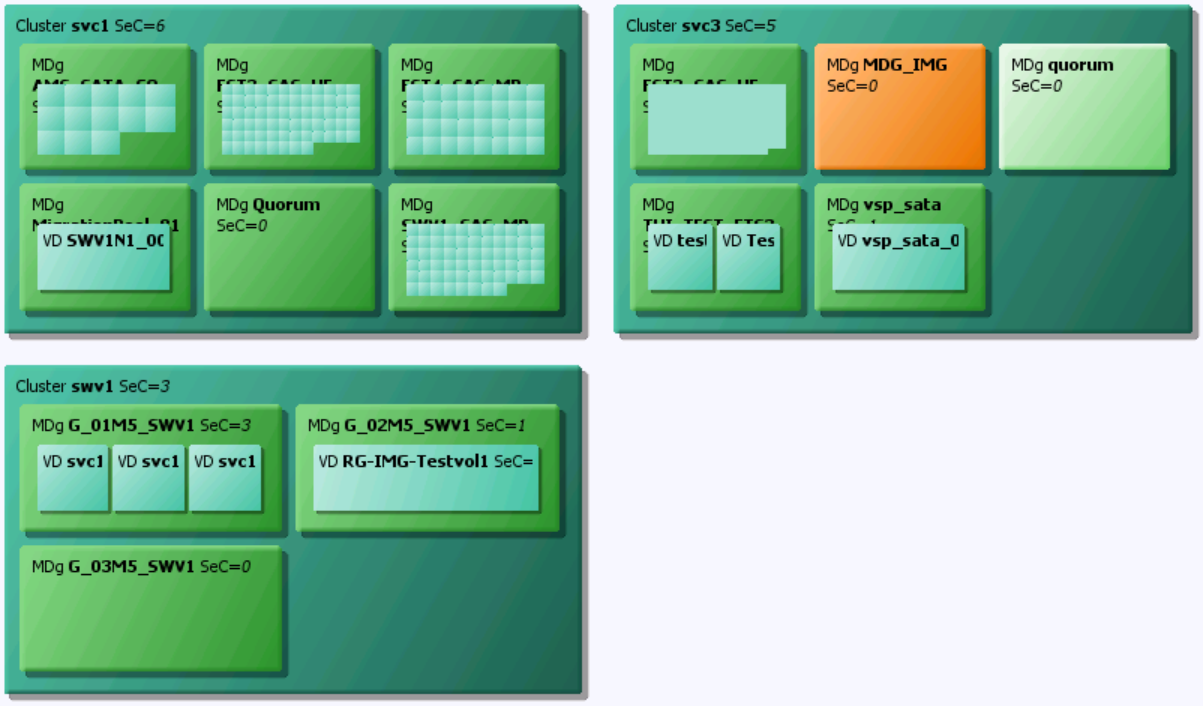

In the picture, the cluster represents the largest, encircling, grouping unit. All areas within the cluster (MDisk group, VDisk, etc.) fill the total area of the cluster-rectangle and reflect the topology of the selected objects of the SVC.

1.5.1.2 Treemap interaction

Interactive features of the treemap like a context menu, drag & drop

- Context menu of the objects in the treemap
- Drag & drop topology editing
  - Allowed drag & drop operations
  - Warning inflicted drag & drop operations
  - Ambiguous drag & drop operations

Context menu of the objects in the treemap

The context menu contains items for the selected object. It is opened by clicking on a particular object with the **right mouse button.**

Tabelle 20 Functions of the treemap´s context menu

	Select subelements (n)	Selects all (n) objects displayed of the next higher level.
	Edit	Modifies attributes of the object, invokes object editors
	Add	Adding new objects with the object editors
	Delete	Deletes object
	Show in list	Calls the tables view of the selected objects in the treemap. They are shown as selected here too.
	Show details	Calls the object detail view of the selected object, also possible by **double click**
	Open SVC CLI	Opens an SSH PuTTY session to the SVC cluster. See: hyperlink auf putty
	Show VDisk performance graph(s)	Calls performance view for linked SVC VDisks
	Show MDisk performance graph(s)	Calls performance view for linked SVC MDisks
	Show MDisk grp. cache perf. graph(s)	Calls performance view for linked SVC MDisk group cache Information
	Show node performance graph(s)	Calls performance view for linked SVC node
	Show node Port performance graph(s)	Calls performance view for linked node ports
	Show node to node performance graph(s)	Calls performance view for linked SVC node to node communication
	Zoom to selection	Centers and zooms the treemap on all selected objects
	Reset zoom (Ctrl+0)	Resets the treemap magnification

Shaded menu items cannot be opened because a precondition is not met.
Example: At a BVQ site, which is associated with no SVC VDisks the menu item "show VDisk performance graph(s)" remains inactive.

Drag & drop topology editing

With the drag & drop function of the treemap relocating SVC VDisk copy-objects between SVC MDisk groups is very easy. Moving allowed objects in the treemap opens the dialog "Cluster CLI Scripts" which collects the available commands in the SVC and is also able to execute them.

It is also possible to move multiple objects (selected by holding down [CTRL] key) with a single draft.

Abbildung 25 Drag & drop operation in the treemap

Allowed drag & drop operations

Drag & drop moves are allowed only for certain types of objects to specific target object types (see table).
Illegal moves are indicated by the lock symbol ⃠ instead of the mouse pointer.

Object type	Target object type
SVC VDisk copy	SVC MDisk group
SVC VDisk	SVC MDisk group
BVQ VDisk group	SVC MDisk group
BVQ application	SVC MDisk group

Also illegal:

- Transfers between SVC MDisk groups with unequal *extent size*.

- Transfers between SVC MDisk groups which are not in the same SVC cluster/V7000.
- Relocations of objects from a source object to the same target object.

Warning inflicted drag & drop operations

Following drag & drop shifts may lead to problematic results and are indicated by a warning symbol ⚠ instead of the mouse pointer.

- Not enough free extents in the target-SVC MDisk group

Ambiguous drag & drop operations

When several objects are affected in a drag & drop operation, which usually should not be on the same target object, an additional dialog appears where the selection can be further limited. Example: A drag & drop operation includes both SVC VDisk copies of a SVC VDisk that would be shifted to the same SVC MDisk group.

Abbildung 26 Drag & drop: ambiguous operations

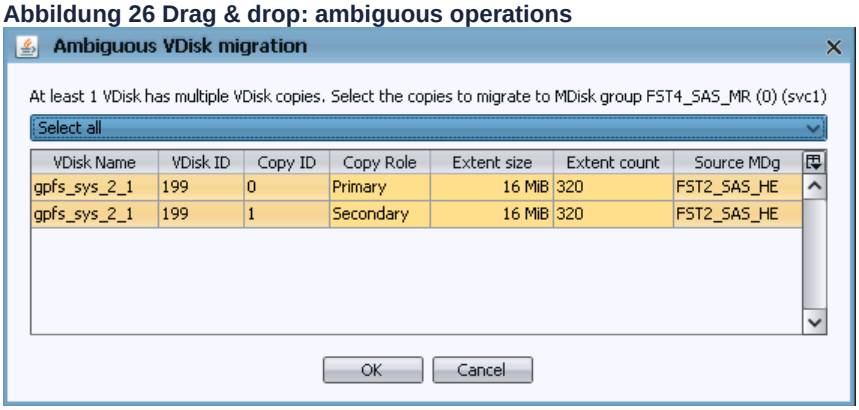
Ambiguous VDisk migration
At least 1 VDisk has multiple VDisk copies. Select the copies to migrate to MDisk group FST4_SAS_MR (0) (svc1)
Select all
VDisk Name
VDisk ID
Copy ID
Copy Role
Extent size
Extent count
Source MDg
gpfs_sys_2_1 199 0 Primary 16 MiB 320 FST2_SAS_HE
gpfs_sys_2_1 199 1 Secondary 16 MiB 320 FST2_SAS_HE
OK
Cancel

1.5.1.3 Treemap navigation

The treemap navigation allows to zoom in and move the image section of the treemap.

Zoom - increase the treemap

With the mouse wheel

To enlarge the tree map, move the mouse wheel

Forward	Scaling up
Backward	Scaling down

The area under the mouse cursor is the center of the zoom. Objects beneath the cursor are usually invisible.

Reset zoom

Abbildung 27 Treemap - reset zoom

Allows to return quickly to the minimum zoom level.

Flash zoom

Double-clicking with the scroll wheel (middle mouse button) maximizes the object-area currently located under the cursor. This allows to zoom quickly into an object level.
For zooming out, a single click on a part of the underlying grouping object-area or the known function "reset zoom" can be used.

Zoom to selection

Abbildung 28 Treemap - zoom to selection

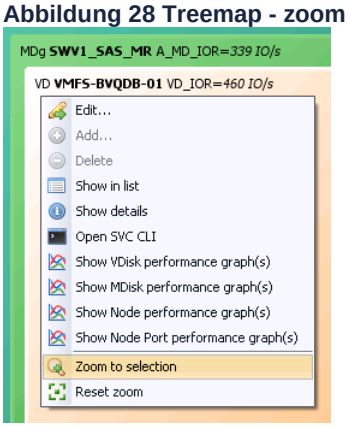

"Zoom to selection" can be called by using the context menu of selected treemap objects. Thus, all currently selected treemap objects are enlarged to the available window area.

Zoom control via keyboard

The treemap zooming can also be operated with the following shortcuts:

Tabelle 21 Shortcuts for the treemap navigation

Shortcut	Function
[Ctrl] [+]	Enlarging treemap
[Ctrl] [-]	Decreasing treemap
[Ctrl] [0]	Minimizing treemap (reset zoom)

Moving of the zoomed-in treemap

Middle mouse button	Detailed moving	Hold and move, mouse cursor "holds" the treemap
[Alt] + left mouse button	Large-scale moving	Hold and move, mouse cursor holds the two scrollbars
Magnifying glass icon	Moving any way	With the first zoom level (zooming via mouse wheel) the magnifying glass icon is displayed. Hold and move, minimap is displayed.

Example: ***minimap of the scroll area***

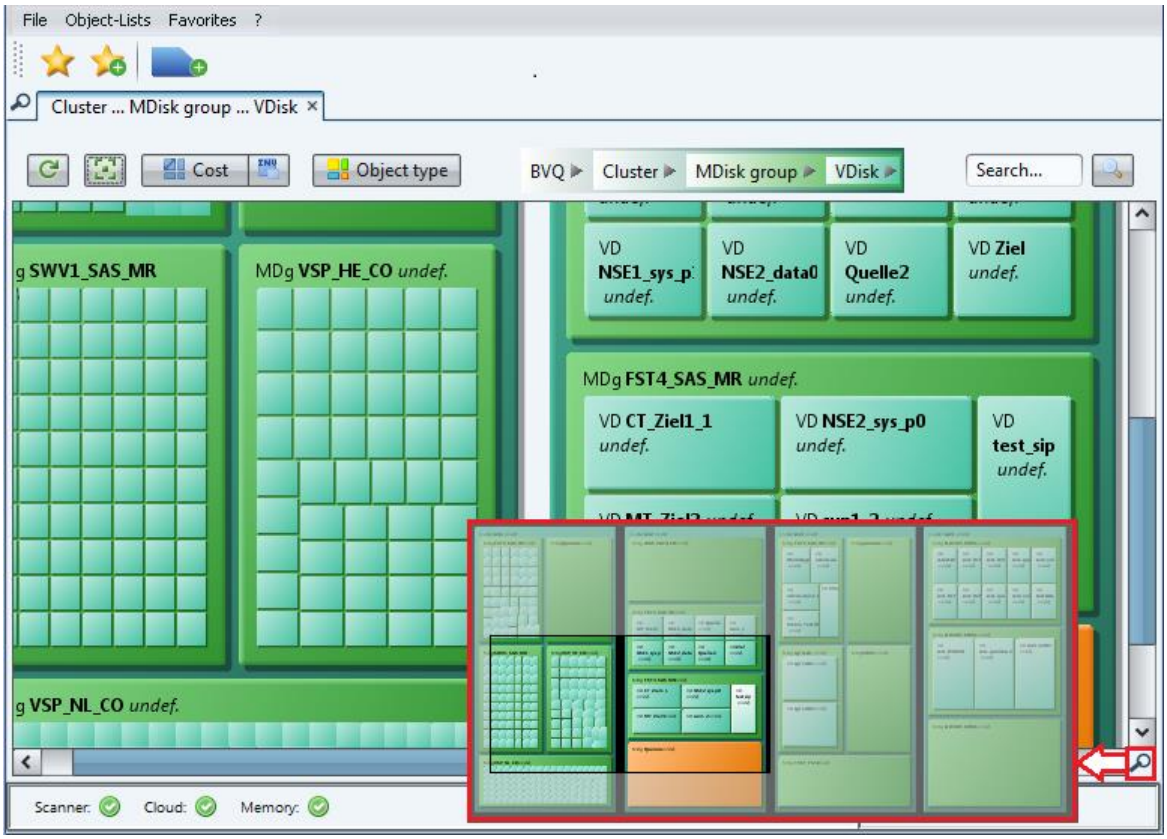

1.5.1.4 Treemap object selection

Selection of objects in the Treemap

In the treemap each object can be chosen (selected) with a mouse click.

To select more than one object of same type, you have to select these objects by holding down the [CTRL] button.
It is not possible to select multiple objects of different types at the same time.

Abbildung 29 Treemap - multiselect with

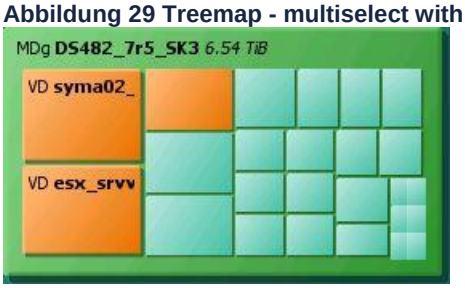

Objects can also be selected with the hit list topology - search function or via the context menu item "select subelements" of an object.

Treemap - select subelements

The context menu of the treemap contains the point of "select subelements". This selects all displayed subelements of the next higher level.

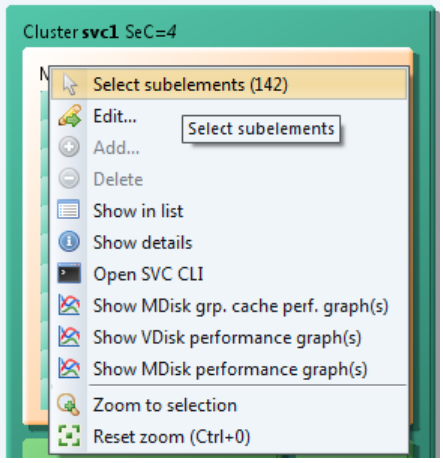

→

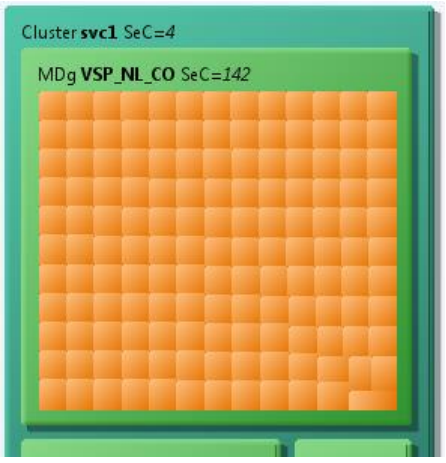

Hereby a large number of objects with common context can be selected easily.

Tabelle 22 Treemap: properties for object selection

Single select	Selects an object with left or right mouse click.
Multi select	Selects multiple objects by mouse click and holding down [CTRL] button.
Subelement select	Selection of displayed subobjects of the selected object with context menu "select subelements".
Search field select	Results of a search query in the search field, are displayed directly in the treemap.
Selection color	Selected objects are marked **orange**.
Selection is retained	Selection is preserved even when path, aspect, coloration and favorites are switched.

1.5.1.5 Treemap tooltips

Treemap tooltips contain essential aspect specific object information.

Move the mouse pointer to the desired object and hold on. Shortly afterwards a tooltip appears with aspect specific information.

For example, in the performance aspect the tooltip shows the performance indicators of the selected integration period.

Or in the capacity aspect information about free and used capacity are displayed.

Abbildung 30 Treemap - tooltip of a VDisk in the aspect

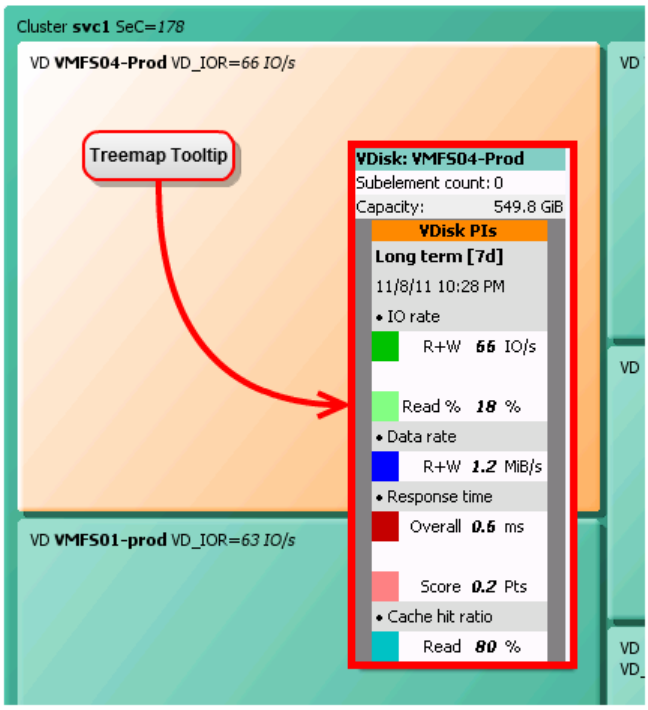

1.5.2 Topology - Object path - Topology navigator

With the topology navigator several BVQ object types are connected to one path, which determines their hierarchical representation within the treemap.

Abbildung 31 Topology navigator

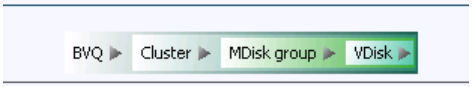

One or more BVQ object types are combined in a path as junction (-points). Each type of object can only occur once in the path. The path is read from left to right and always begins with the root node "BVQ" and ends, not visible, with the last node "nothing".

The possible relationships between object types are predefined in BVQ - combinations without reasonable meaning are filtered out and are not selectable.

1.5.2.1 Navigation

By clicking a node in the path a new drop down menu is opened which allows to insert further nodes after or replace the selected node in the path:

Abbildung 32 Topology navigator - menu for path intersection

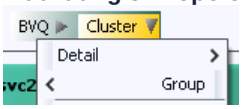

A selection of "BVQ" sets the root node to the beginning of the path. The selection of "-nothing-" terminates the path. These two special nodes are used to cut a path in front, behind or directly at an selected node. It is not necessary to select them on the end of a selection range. A single click on the last or first desired object type is sufficient to completely specify the path.

Detail menu

Extends the path to the right. Other nodes are added behind the selected one. Nodes following the selected one are discarded.

Abbildung 33 Topology navigator - Detail menu

Group menu

Extends the path to the left. New nodes are added in front of the selected one. Nodes prior to the selected one are discarded.

Abbildung 34 Topology navigator - group menu

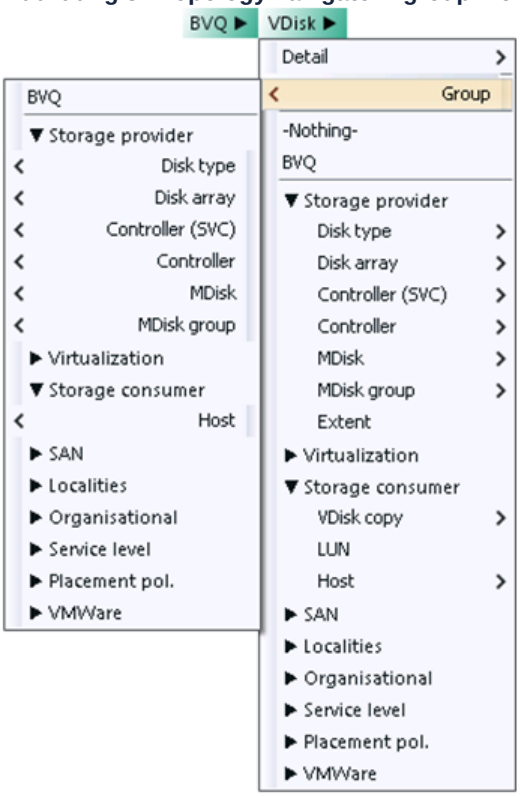

Exchange menu

The object types that can be used in place of the selected node appear below the dividing line after detail / group.

Abbildung 35 Topology navigator - Exchange menu

1.5.2.2 Object type categories

Object types are divided into categories, which can be opened and closed within the entire menu tree.

To prevent unnecessary opening/closing operations, the topology navigator saves the last "drop"-status (open / closed) of each category and applies it to new menus. The topology navigator tries to self-adapt to the user situation.

If a topology tab is initialized from the favorite menu, all categories containing object types in the new path are opened.

Hint: It is convenient to open all categories required by the current application at the first path selection with BVQ → Detail. The rest should be closed or remain. Resulting in only few menu items visible simultaneously, the navigation is simplified.

Abbildung 36 Topology navigator - Object type categories

1.5.2.3 Keyboard control

The path navigator is fully controllable by keyboard

Button	Function
[Alt+l]	Enables the path navigator
[→]	Branches into the submenu
[←]	Back to top menu
[↑] / [↓]	Menu selection marker up/down
[Return]	Activates menu item * Category opens/closes * Object type is selected
[ESC]	Leaves the menu

1.5.2.4 Topology navigator und favorites

Topology hierarchies are stored and then opened by using favorites. Some predefined favorites are provided with the BVQ installation. (see Favorites and tabs)

1.5.2.5 Practice hint: ad hoc analyses

Many analyses are made ad hoc. Most of the times only very flat hierarchies are used. This results in a quick and clear result. In these cases, the settings of the topology navigators are usually not done with favorites, but made by hand.

Two examples:

- Example1 - overview of all the connected servers and their current performance parameters.
  The object path is set to hosts, the aspect to capacity-> total - the server with the highest capacity requirements is also on top in the treemap.
  To see the VDisks on the server, the object VDisks is added to the host in the topology navigator.
  If needed, a performance analysis for each displayed object can now be done with the right mouse button.

- Example2 - which average data processing rates are used by the volumes of the SAP installation?
  Set VDisks in the topology navigator and choose aspect perf. long-> data rate R + W.
  Set a meaningful name-fragment for the SAP discs in the object search. (see Topology-view)
  The found SAP disks are highlighted in orange - details are made visible via zoom or the pop-up window.

1.5.3 Topology - aspects

The aspects of the topology view allow to refocus the display according to certain priorities.

The main function of the aspects is to specify the size factor for the displayed area size in the treemap.
Depending on the aspect´s selection, the objects with the largest capacity, load, cost, ... are also represented as the largest areas.

The following aspect categories are available:

Tabelle 23 Topology - aspect selection

Aspect topology	The aspect *topology* is focused on the representation of the objects´ hierarchical order.
Aspect capacity	The aspect *c*apacity is targeted to management of storage capacities.
Aspect performance	The *performance* aspect is focused on the analysis of IO load/performance and visualizes predicted load- / performance indicators.
Aspect cost & price	*Price and cost* aspects are focused on the representation of specific financial accounting variables.
Inv	Object type in the top level of the treemap is displayed inverse proportionally to the selected size factor (inverted).

The selection of an aspect affects several functional units of the display:

- Treemap
  - Metric by which the **surface area** of the objects is aligned (proportionality factor - PF)
  - Labeling of the object area
  - Content of the object tooltips
- Object detail view
  - Selection of the chapters in the object detail view.

⚠Only absolute values are shown in the treemap. Actually displayed sub-objects in a surrounding object-area are irrelevant (no subelement-sizing).

⚠Not all objects have a unique hierarchical mapping to each other. Example: performance of a MDisk group is calculated by using MDisks and not VDisks. Total VDisk IO within the MDg is rarely equal to the MGg IO.

1.5.3.1 Aspect topology

The aspect *topology* is focused on the representation of the objects´ hierarchical order.

All objects with the same path level are displayed in the treemap by equally sized squares.

The presentation focuses on the questions:

- How are objects arranged in the path hierarchy?
- How many objects belong to a group?
- What is the topological relation of the objects?

Particularly important here is the following parameter:

SeC - Subelement count	Displays the number of child-objects of a grouping object in the current treemap.

The sub element count is displayed in the caption area and in the objects´ tooltip.

Abbildung 37 Subelement count

1.5.3.2 Aspect capacity

The aspect capacity is targeted to management of storage capacities.

Abbildung 38 Topology aspect - capacity

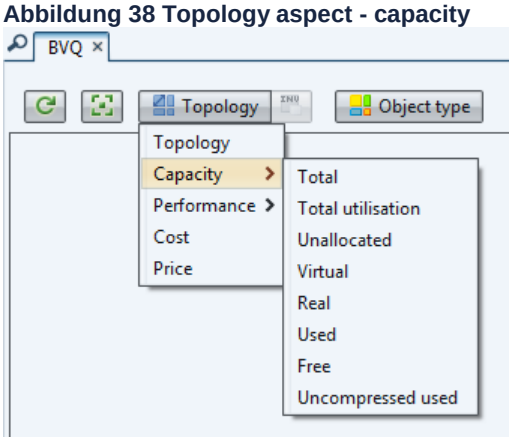

The aspect capacity paints the object sizes of the treemap in relation to each objects´ storage capacities metric.

It is a collective term for the following sub-aspects:

Tabelle 24 Aspect capacity - sub-aspects

Aspect	Area size is proportional to
Total	TOTAL capacity
Total utilisation	Total storage capacity combined with unused capacity
Unallocated	UNALLOCATED capacity
Virtual	"Volume VIRTUAL capacity" *With Thin Provisioning: the capacity visible to the host*
Real	"Volume REAL capacity" *With Thin Provisioning: the capacity that is currently used by the pool*
Used	"Volume USED capacity"
Free	"Volume FREE capacity"
Uncompressed used	"Volume UNCOMPRESSED_USED capacity"

Application example - total utilisation

In this aspect, the displayed storage provider objects´ areas are divided into two areas:

- **Top** - blank (free) capacity
- **Bottom** - used capacity - where the storage consumer objects are displayed

This aspect provides extended representations for the following paths. For all other paths, it shows the same view as in the aspect *capacity total*.

Tabelle 25 Topology aspect capacity - total utilisation: meaningful object type paths

MDisk group	►	VDisk
MDisk group	►	VDisk copy
MDisk group	►	Extent
MDisk	►	VDisk
MDisk	►	VDisk copy
MDisk	►	Extent

⚠ The path combinations [MDISK group, MDISK] ► [VDisk] may lead to misinterpretations, if the VDisk is mirrored by *VDisk mirror*. Then the area of the VDisk corresponds to the "real capacity" sum of both copies.

Abbildung 39 Topology aspect capacity - total utilisation

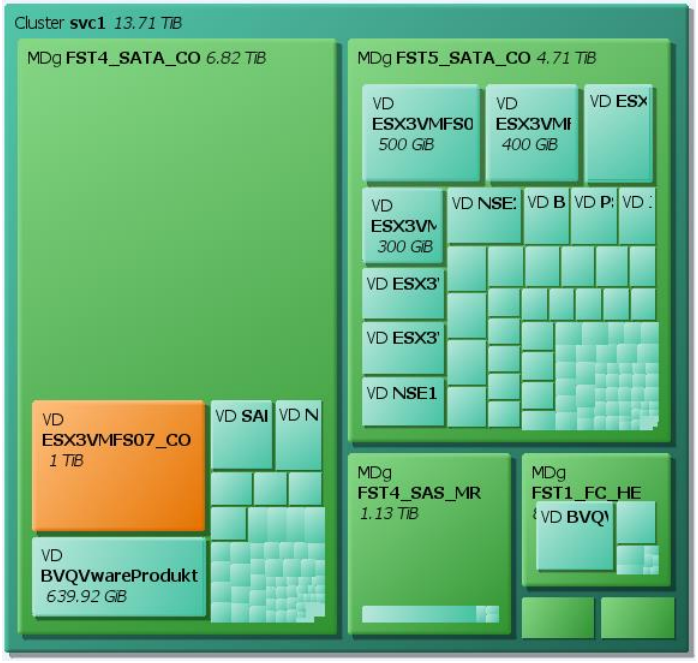

It's clear to see that all VDisk areas, from whatever MDISK group, are now comparable to each other.
Just looking at the chart it becomes obvious, the VDisk "ESX3VMFS0" has enough free space in the MDISK group "FST4_SATA_CO".
The displacement of capacities by drag and drop to optimize capacity allocation is particularly simplified just by this aspect.

1.5.3.3 Aspect performance

The *performance* aspect is focused on the analysis of IO load/performance and visualizes predicted load- / performance indicators.

The objects´ area size is displayed proportionally to the selected performance indicator.

The selection menu for the aspect performance looks like:

Abbildung 40 Topology aspect - performance

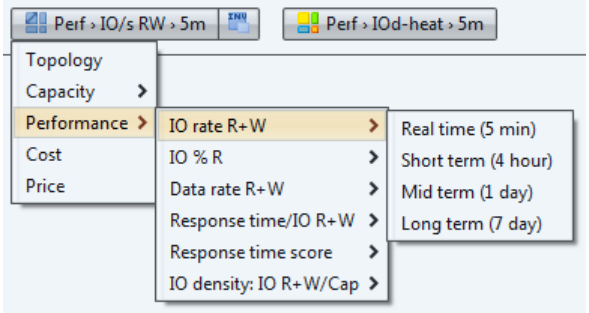

The selectable periods depend on the BVQ license!

Load values of the object types SVC VDisk, SVC MDisk and SVC node are read from the BVQ DB in evaluation runs and condensed to performance indicators.
More information about the selectable periods and metrics can be found in the chapter performance indicators.

IO density:

The metric *IO density : IO R+W/cap* is different. It also serves as basis for the color scheme IO density heat. Principles for the calculation can be found there.
In contrast to the respective color scheme, no filtering is applied in this aspect. Thus, objects with very small used capacity appear to be very large, despite smallest IO (typical for e.g. cluster Heartbeat volumes or newly created thin provisioned volumes).

The visualization of load indicators in the treemap of the topology view is complemented by the analysis of the historical load profile in the performance view.

An important part of BVQ is the support of the performance analysis.
Typical goals of such analyzes are:

- The localization of performance problems
- Finding the volume, which is the potential source of the excessive load. Which volumes just suffer from the side effects?
- Optimizing the utilization or costs by redistributing to cheaper storage.
- Finding the storage consumers with the highest load (consumer perspective).
- Finding of resources (storage providers) with the most free load-reserves (provider perspective).

Here is an example for the representation of the topology in a performance aspect:

Abbildung 41 Topology aspect performance - view mid-term - I/O rate R+W

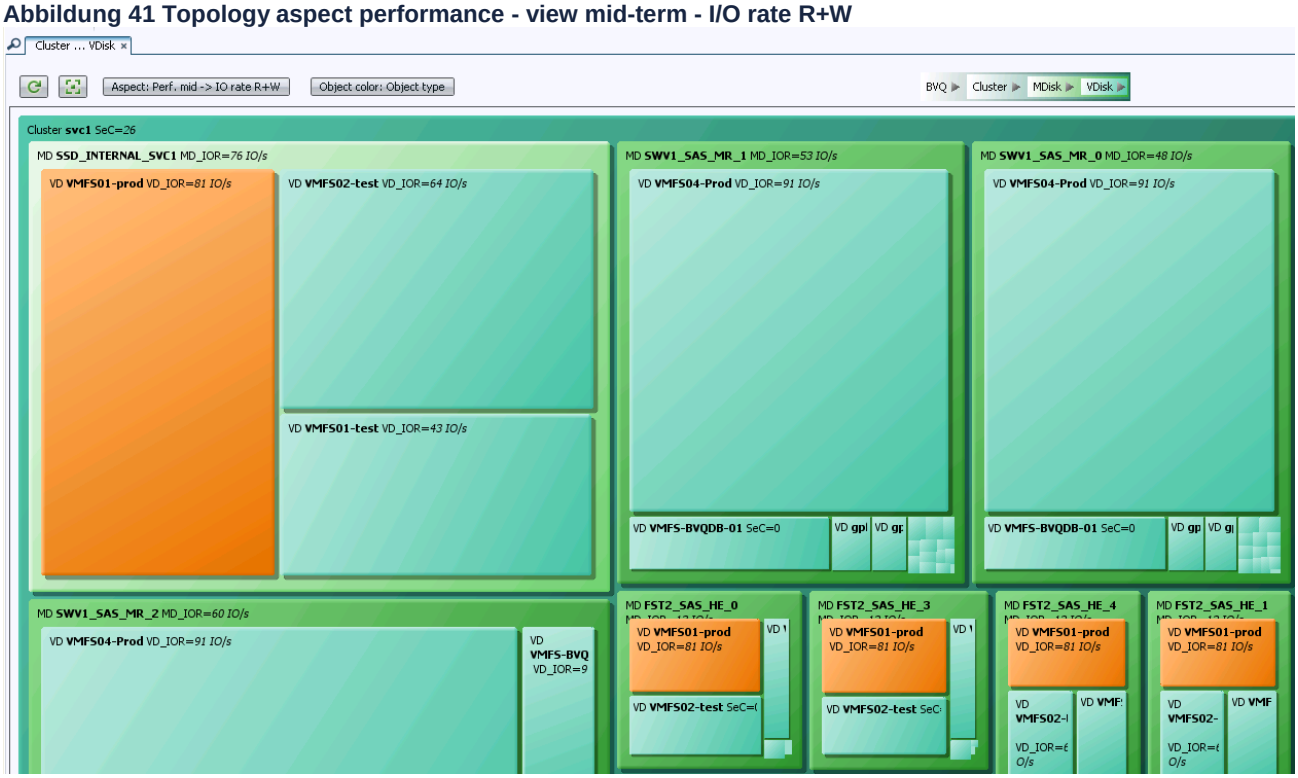

1.5.3.4 Aspect cost & price

Price and cost aspects are focused on the representation of specific financial accounting variables.

Aspect	Multiplier	Purpose
Cost	MDisk group cost ratio	Shows costs based on the total capacity of storage providers or the capacity occupied by storage consuming objects.
Price	VDisk group current SLA price ratio	Displays sales price based on the virtual capacity of storage consuming objects.

Aspect cost

The cost aspect allows the examination of the costs incurring by storage providers. Before cost aspects can be used, the storage pools (MDISK groups) have to be configured with the factor "cost ratio" .

⚠ Depending on the BVQ edition not all cost aspects are available.

Abbildung 42 Topologie Aspekt Cost - Kostenbetrachtung

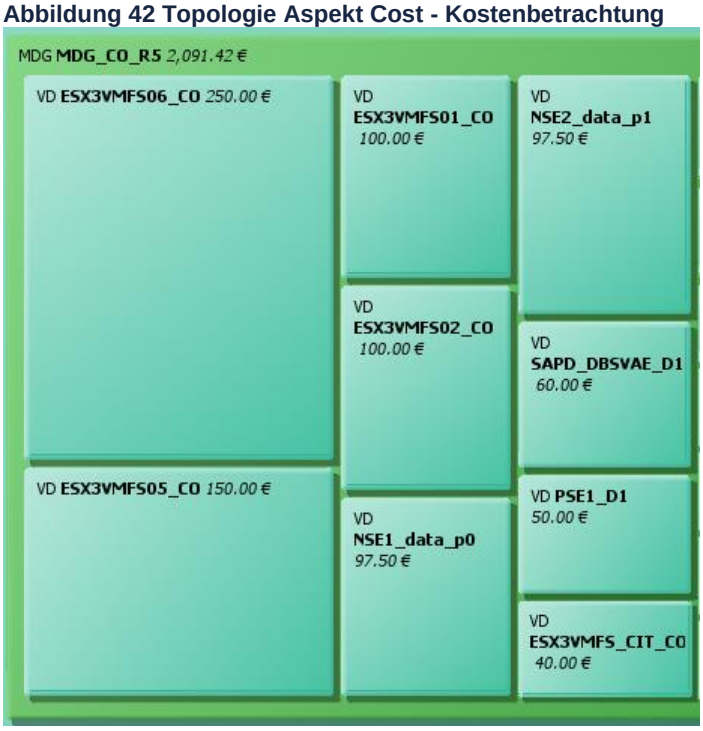
MDG MDG_CO_R5 2,091.42 €
VD ESX3VMFS06_CO 250.00 €
VD ESX3VMFS01_CO 100.00 €
VD NSE2_data_p1 97.50 €
VD ESX3VMFS02_CO 100.00 €
VD SAPD_DBSVAE_D1 60.00 €
VD ESX3VMFS05_CO 150.00 €
VD NSE1_data_p0 97.50 €
VD PSE1_D1 50.00 €
VD ESX3VMFS_CIT_CO 40.00 €

Aspect price

The price aspect allows the examination of the prices which are calculated by the storage capacity for every user (consumer/customer). Before price aspects can be used, the SLAs have to be configured with the factor "price ratio" .

⚠ Depending on the BVQ edition not all price aspects are available.

Abbildung 43 Topology aspect price - price analysis

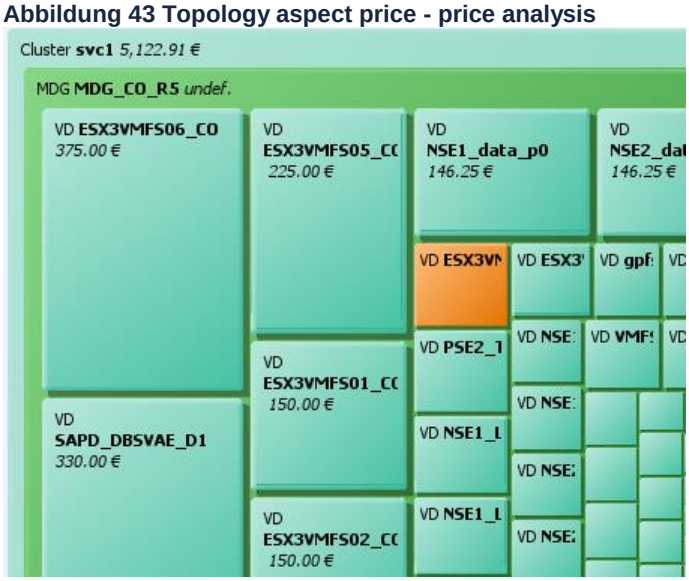

1.5.4 Topology - color schemes

Via the color schemes the objects´ areas in the treemap are colored, based on selectable criteria.

To enable a quick and meaningful examination, different BVQ color schemes can be used, coloring the objects according to certain criteria. Partially, these object colors can be designed freely. The dynamic coloring of the treemap-objects simplifies the conditon-assessment of the objects like qualification, service level and compliance.

A typical example: all not SLA-compliant objects are shown in red, while the objects which satisfy all SLA-conditions are stained in green.

1.5.4.1 Main menu

The main menu for the color scheme looks like:

Abbildung 44 Topology - color scheme: main menu

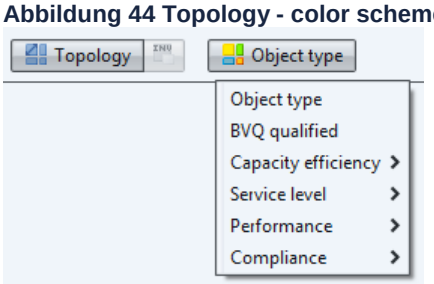

The following color schemes can be selected:

- Color scheme object type — The object coloring according to the scheme "object type" colors the objects in the treemap depending on the BVQ integrated standard color of the respective object types.

- Color scheme BVQ qualified — The object staining according to the scheme "BVQ qualified" colors the treemap objects depending on the state of the BVQ qualification process.
- Color scheme total capacity efficiency — The object staining according to the scheme "total capacity efficiency" colors the objects based on actual saving of total capacity.
- Color scheme space efficiency — The object coloring scheme "space efficiency" paints the objects of the treemap based on the data reduction by Thin Provisioning.
- BVQ color scheme compression efficiency — The color scheme "compression efficiency" colors objects according to the the data-reduction efficiency achieved with compression.
- Color scheme compression rate — The object coloring according to the scheme "compression rate", colors the objects based on the compression rate.
- Color scheme data and storage class — The object staining according to the scheme "data and storage class" colors the objects in the treemap according to the custom colors of the BVQ storage class or BVQ data class.
- Color scheme DC2SC placement rating — The object staining, according to the scheme "DC2SC placement rating", colors the objects of a treemap based on the placement review of the associated BVQ DC to SC placement ratings.
- Color scheme IO density heat — The object coloring according to the scheme "IO density heat" paints the objects of the treemap based on an assessment by comparing the configurable estimated IO density of the object type BVQ disk array.
- Color scheme response time compliance — The object coloring according to the scheme "response time compliance" colors the objects of a treemap based on an assessment of the configurable design response time from the object type SVC MDisk group.
- Color scheme overall compliance — The object staining according to the scheme "overall compliance" colorizes the objects of the treemap according to the combined status of all compliance conditions.
- Color scheme data class compliance — The object color scheme "data class compliance" colors the objects of the treemap according to the SLA conformity, defined in the BVQ DC to SC placement rating.
- Color scheme mobility zone compliance — The object coloring according to the scheme "mobility zone compliance" colors the objects of the treemap on the compliance-state with requirements defined in the object type BVQ mobility zone.
- Color scheme VDg relation compliance — The object coloring according to the scheme "VDG relation compliance" colors the objects of the treemap on the state of requirements´ compliance defined by the object type BVQ VDisk group relation.
- Object type specific color schemes — Some object types have a specially adapted color scheme for a better representation of their specific properties.

1.5.4.2 Color scheme group: capacity efficiency

With the object colorings of the scheme-group *capacity efficiency*, the compliance of Thin Provisioning (space efficency) and the compression thresholds can be tested in the treemap.

Abbildung 45 Topology - color scheme: submenu

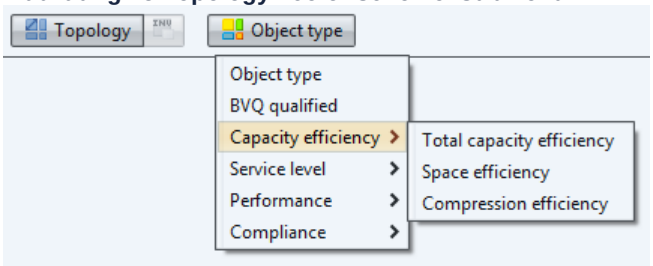

The following color schemes belong to this group:

- Color scheme total capacity efficiency — The object staining according to the scheme "total capacity efficiency" colors the objects based on actual saving of total capacity.
- Color scheme space efficiency — The object coloring scheme "space efficiency" paints the objects of the treemap based on the data reduction by Thin Provisioning.
- BVQ color scheme compression efficiency — The color scheme "compression efficiency" colors objects according to the the data-reduction efficiency achieved with compression.

1.5.4.3 Color scheme group: service level

With the object colorings of the scheme-group *service level,* the compliance of service level rules configured by BVQ, can be checked in the treemap.

Abbildung 46 Topology - color scheme: submenu

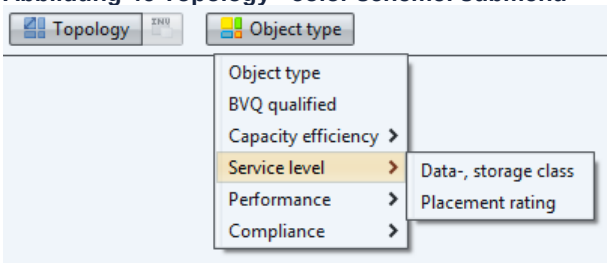

The following color schemes belong to this group:

- Color scheme data and storage class — The object staining according to the scheme "data and storage class" colors the objects in the treemap according to the custom colors of the BVQ storage class or BVQ data class.
- Color scheme DC2SC placement rating — The object staining, according to the scheme "DC2SC placement rating", colors the objects of a treemap based on the placement review of the associated BVQ DC to SC placement ratings.

1.5.4.4 Color scheme: performance

With the object colorings of the scheme-group *performance* the compliance ot the BVQ configured load-threshold values can be checked in the treemap.

Abbildung 47 Topology - color scheme: sub menu

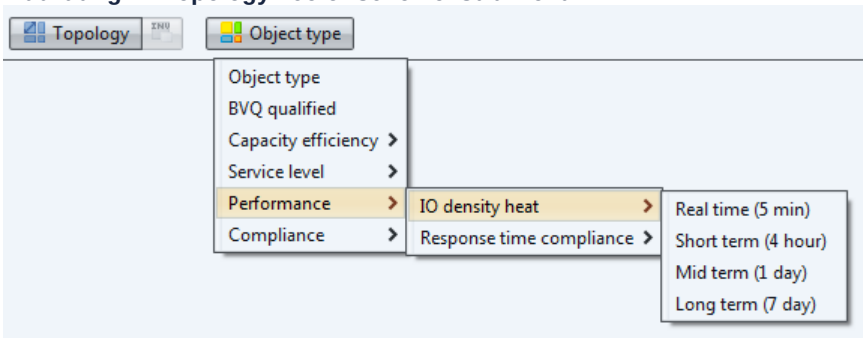

The following color schemes belong to this group:

- Color scheme IO density heat — The object coloring according to the scheme "IO density heat" paints the objects of the treemap based on an assessment by comparing the configurable *estimated IO density* of the object type BVQ disk array.
- Color scheme response time compliance — The object coloring according to the scheme "response time compliance" colors the objects of a treemap based on an assessment of the configurable *design response time* from the object type SVC MDisk group.

1.5.4.5 Color scheme group: compliance

With the object stainings of the object scheme-group *compliance* the compliance of BVQ configured placement ruels can be checked in the treemap.

Abbildung 48 Topology - color scheme: submenu

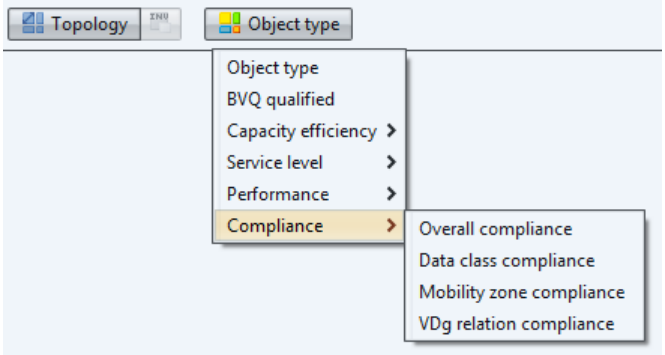

The following color schemes belong to this group:

- Color scheme overall compliance — The object staining according to the scheme "overall compliance" colorizes the objects of the treemap according to the combined status of all compliance conditions.
- Color scheme data class compliance — The object color scheme "data class compliance" colors the objects of the treemap according to the SLA conformity, defined in the BVQ DC to SC placement rating.
- Color scheme mobility zone compliance — The object coloring according to the scheme "mobility zone compliance" colors the objects of the treemap on the compliance-state with requirements defined in the object type BVQ mobility zone.

- Color scheme VDg relation compliance — The object coloring according to the scheme "VDG relation compliance" colors the objects of the treemap on the state of requirements´ compliance defined by the object type BVQ VDisk group relation.

1.5.4.6 Object type specific color schemes

Some object types have a specially adapted color scheme for a better representation of their specific properties. See object type specific color schemes

1.5.4.7 Color scheme object type

The object coloring according to the scheme "object type" colors the objects in the treemap depending on the BVQ integrated standard color of the respective object types.

BVQ standard colors for each object type category:

Tabelle 26 BVQ standard colors for object types

Color	Object type - category
Green	Storage provider
Turquoise	SVC virtualization
Blue	Storage consumer
Lilac	BVQ - VMware virtualization

Some object types have a specially adapted color scheme for a better representation of their specific properties. Please see object type specific color schemes

1.5.4.8 Color scheme BVQ qualified

The object staining according to the scheme "BVQ qualified" colors the treemap objects depending on the state of the BVQ qualification process.

The colors have the following significance:

Tabelle 27 Topology-color scheme

Color	Symbol	Condition	Description
Green		Complete	All necessary additional BVQ-information is completely configured
Red		Incomplete	Some details are still missing
Object type		Inapplicable	Type BVQ-qualified state is not defined for this object

Rules for the state BVQ qualified:

Tabelle 28 BVQ qualified rule set for each object type

Object type	Completely qualified when grouped into
BVQ room	BVQ site
BVQ application group	BVQ cost center
BVQ application	BVQ application group
BVQ VDisk group	BVQ application BVQ SLA
SVC VDisk	BVQ VDisk group
SVC controller	BVQ controller
BVQ controller	BVQ room
SVC node	BVQ room
SVC host	BVQ room
BVQ disk array	BVQ controller
SVC MDisk	BVQ disk array
SVC MDisk group	BVQ storage class
BVQ storage class	BVQ DC to SC placement rating for every BVQ data class
BVQ data class	BVQ DC to SC placement rating for every BVQ storage class
BVQ SLA	BVQ data class
BVQ mobility zone	SVC MDisk group
BVQ VDisk group relation	BVQ VDisk group

1.5.4.9 Color scheme total capacity efficiency

The object staining according to the scheme "total capacity efficiency" colors the objects based on actual saving of total capacity.

Tabelle 29 Color scheme capacity efficiency - meaning of the colors

Color	Efficiency	Value range	Meaning
Red	Poor	0 - 24%	Poor compression rate or low data reduction by Thin

			Provisioning
Yellow	Average	25 - 49%	Average compression rate or low data reduction by Thin Provisioning
Green	Good	50 - 79%	Average compression rate or low data reduction by Thin Provisioning
Blue	Excellent	80 - 100%	Excellent compression rate or highest data reduction by Thin Provisioning
Faded	Not Applicable	--	Affected object type, but object doesn´t have the right data-reduction type for the selected color
Object type	Not Applicable	--	Object type is not affected

Abbildung 49 Color scheme total capacity efficiency - treemap & show details screenshot

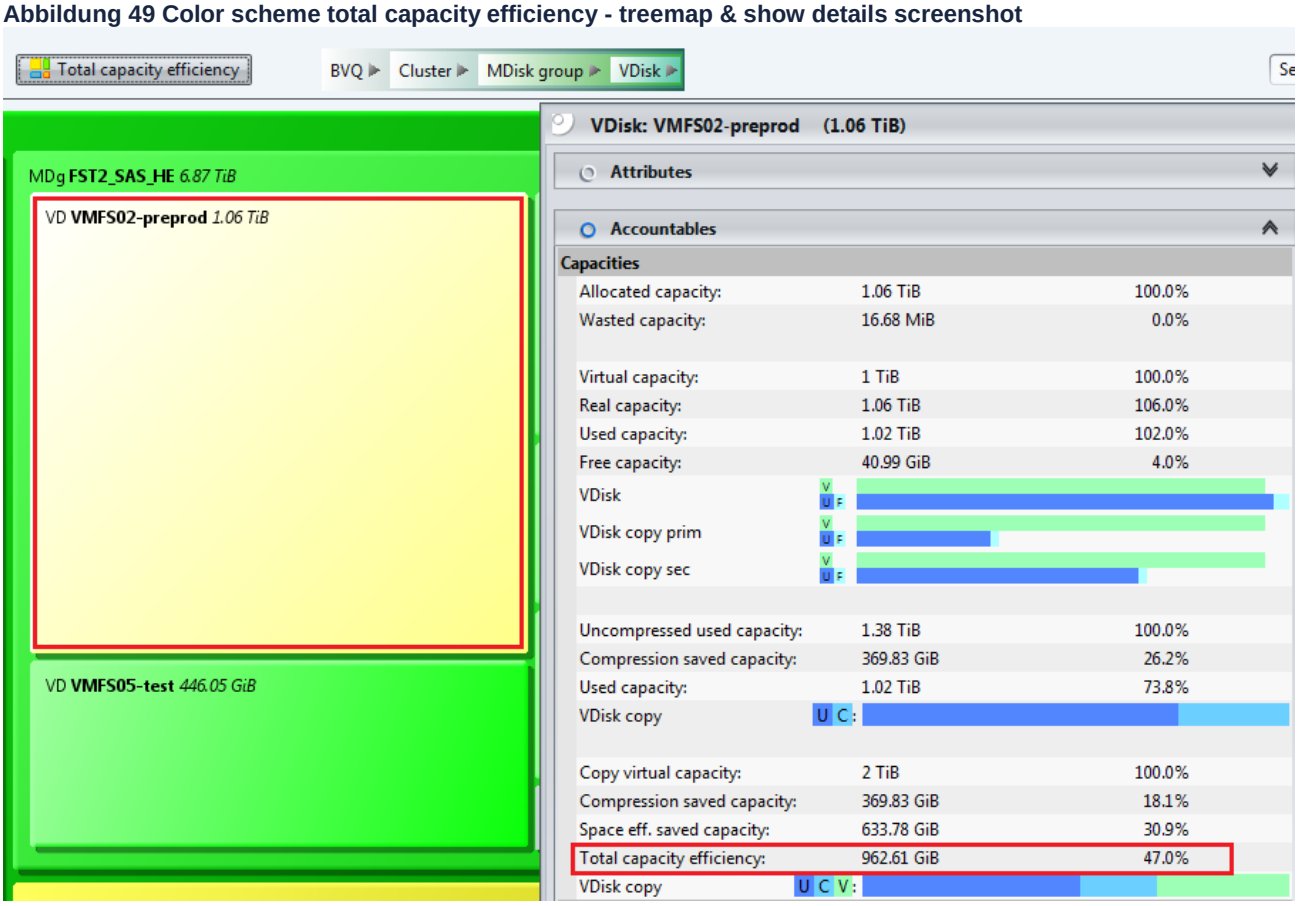

1.5.4.10 Color scheme space efficiency

The object coloring scheme "space efficiency" paints the objects of the treemap based on the data reduction by Thin Provisioning.

Tabelle 30 Color scheme space efficiency - meaning of the colors

Color	Efficiency	Value range	Meaning
Red	Bad	0 - 24%	Low data reduction by thin provisioning
Yellow	Average	25 - 49%	Average data reduction by thin provisioning
Green	Good	50 - 79%	High data reduction by thin provisioning
Blue	Excellent	80 - 100%	Very high data reduction by thin provisioning
Faded	Not space eff.	--	Object is not configured for space efficiency
Object type	Not applicable	--	Object type is not affected

Abbildung 50 Color scheme space efficiency - treemap & show details screenshot

1.5.4.11 BVQ color scheme compression efficiency

The color scheme "compression efficiency" colors objects according to the the data-reduction efficiency achieved with compression.

Tabelle 31 Color scheme compression efficiency - meaning of the colors

Color	Efficiency	Value margin	Meaning
Red	Poor	0 - 24%	Insufficient data-reduction with compression
Yellow	Average	25 - 49%	Average data-reduction with compression
Green	Good	50 - 79%	High data-reduction with compression
Blue	Excellent	80 - 100%	Very high data-reduction with compression
Eclipsed	No compression	--	Object is not configured for compression
Object Type	Not applicable	--	Object type is not involved

Abbildung 51 Color scheme compression efficiency - treemap & show details

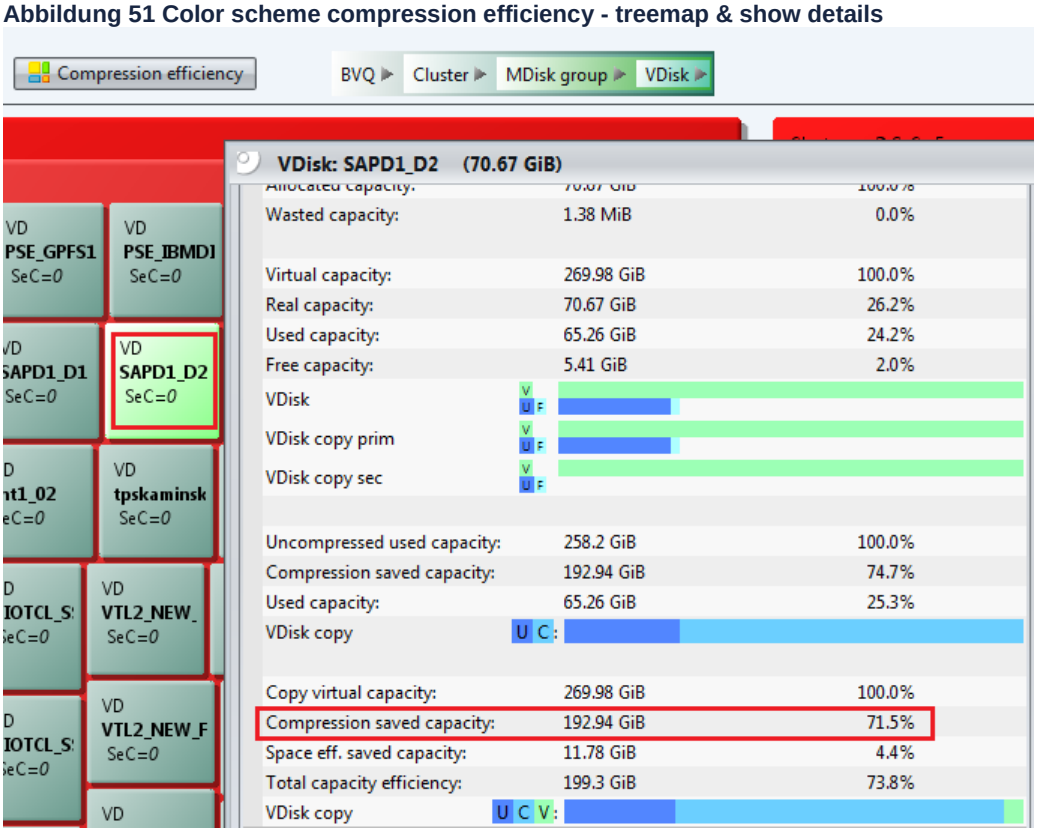

1.5.4.12 Color scheme compression rate

The object coloring according to the scheme "compression rate", colors the objects based on the compression rate.

Tabelle 32 Color scheme compression rate - meaning of the colors

Color	Efficiency	Value range	Meaning
Red	Bad	0 - 24%	Low compression rate
Yellow	Average	25 - 49%	Average compression rate
Green	Good	50 - 79%	High compression rate
Blue	Excellent	80 - 100%	Very high compression rate
Faded	No compression	--	Object is not configured for compression
Object type	Not applicable	--	Object type is not affected

1.5.4.13 Color scheme data and storage class

The object staining according to the scheme "data and storage class" colors the objects in the treemap according to the custom colors of the BVQ storage class or BVQ data class.

Depending on an object´s belonging, it is colored in the same color like the storage or data class:

Tabelle 33 Topology color scheme

Color	Condition
Color BVQ storage class	Objects with configured storage class
Color BVQ data class	Objects with configured data class
Object type	Ambiguously assigned objects (e.g., VDisk copies of a Volume are stored in different DCs.),or object types are not affected by this color scheme (e.g. ports)
Object type faded	No configured storage or data class is available for this object

1.5.4.14 Color scheme DC2SC placement rating

The object staining, according to the scheme "DC2SC placement rating", colors the objects of a treemap based on the placement review of the associated BVQ DC to SC placement ratings.

In the treemap is indicated whether the respective objects, configured due to BVQ specifications, are in a suitable storage space.

Possible colors of this color scheme are:

Tabelle 34 Topology color scheme

Color	Symbol	Condition	Description
Green		Optimal	Placement review is optimal
Yellow		Acceptable	Placement assessment is acceptable or too risky or too expensive
Red		Unacceptable	Placement review is unacceptable (perhaps not SLA compliant)
Faded		Undefined	Requirements for classification are not sufficiently configured or are not convenient to object
Object type		Inapplicable	Object type is not affected

1.5.4.15 Color scheme IO density heat

The object coloring according to the scheme "IO density heat" paints the objects of the treemap based on an assessment by comparing the configurable *estimated IO density* of the object type BVQ disk array.

Determination of the threshold value per storage pool (SVC MDisk group):

For each BVQ disk array the so-called "estimated IO density", a threshold value for the IO density, can be defined.
It is automatically calculated on the basis of the underlying BVQ disk type and the RAID level of the BVQ disk arrays.

The IO density threshold for correlating SVC MDisk groups is also calculated by BVQ based on this threshold.

Calculation of the measured IO density:

The currently measured IO density is based on the performance indicators "IO rate" (IOR) minus the "read cache hit rate %R" (RCH%), available in the known integration intervals. This IO rate is approximately equal to the storage system IO rate (also neglecting the write cache hits on dirty data).
The result is further divided by the storage capacity on which the IO can potentially be distributed (IOCap). Depending on the object´s type a matching metric from the storage capacities is used.

- Storage providers: ALLOCATABLE capacity
- Storage consumers: USED capacity

Entire formula:
IOdensity = (IOR - IOR*RCH%) / IOCap

For VDisk copies and their groupings the read and write IO is attributed to the primary side. The secondary copies only get the write IO.

The determined IO density can also be displayed in the aspect performance.

Comparison of threshold values:
This color scheme compares the IO density threshold of the MDISK groups with the currently measured IO density of correlated objects and stains them as follows:

Tabelle 35 Color scheme IO density heat - meaning of the colors

Color	Status	Meaning
Red	Hot	Current IO density is above the threshold
Yellow	Warm	Current IO density is above 80% of the threshold
Green	Normal	Current IO density is below 80% of the threshold
Blue	Cold	Current IO density is below 20% of the threshold OR IOrate < 0,5 IO/s
White	Filtered	Capacity < 1GiB AND 0,5IO/s < IOrate < 10IO/s
Faded	Undefined	No thresholds are defined for the object OR IOrate = undef.
Object type	Not Applicable	Object (-type) is not affected

Status-aggregation:

- Higher grouping **storage consuming** object types, like VDisk, VDisk group, Data class ... aggregate the status based on the object type *VDisk copy*.
  For the "primary copy" only the read & write IOs are counting and for the further copies merely the write IOs .
- Higher grouping **storage providing** object types, like MDisk group, storage class ... aggregate the status based on the object type *MDisk*.

Here, the "hottest" status determines the aggregation result.

IO density in the object detail view

In the object detail view under "performance" the chapter "IO density" can be found:

Abbildung 52 Object detail-view: extract performance -> IO density

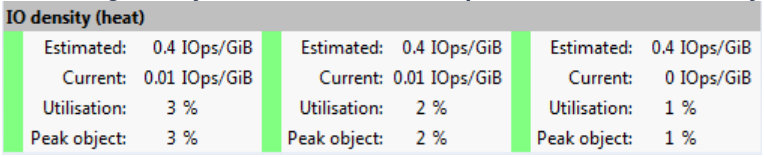

Estimated I/O	The estimated maximum I/O density of the object, based on the *estimated IO*

density	*density* of the object BVQ disk array.
Current I/O density	The objetc´s measured I/O density. Hence the object´s area size is derived in the aspect *perfomance -> IO density*.
Utilisation	The object´s measured I/O density based on the *estimated I/O density*
Peak object	The measured I/O density of the most used sub-object (only for aggregating object types) based on its *estimated I/O density.* Hence the object staining is derived in this color scheme.

How-to: usage of coloring strategy IO density heat

1. Define the *estimated IO rate* in the BVQ disk type and *the estimated IO density* in the BVQ disk array and connect the arrays with the correlating SVC MDisks.
2. Chose the following *predefined favorite*: "**Analysis performance -> IO density heatmap: cluster ... MDg ... VDisk copy**"
3. The result is the following visualization of a heatmap in the treemap

Abbildung 53 Color scheme IO density heat - screenshot heatmap

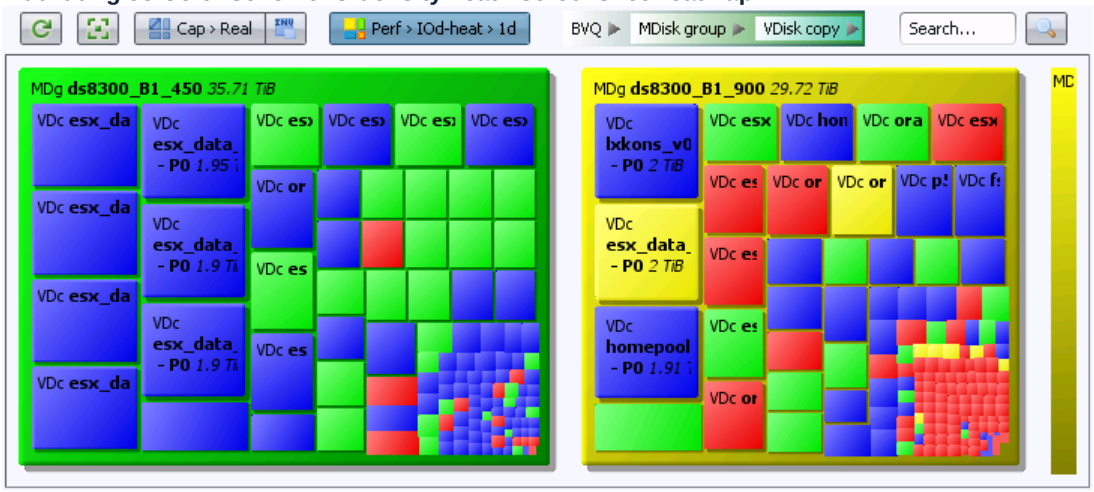

1.5.4.16 Color scheme response time compliance

The object coloring according to the scheme "response time compliance" colors the objects of a treemap based on an assessment of the configurable *design response time* from the object type SVC MDisk group.

For each MDisk group a threshold for the response time can be set, the so-called "design response time."

The measured *average response time* is based on the performance indicator "response time R/W", available in the known intervals.

This color scheme compares the response time threshold of MDisk groups with the currently measured response time of correlated objects and stains them as follows:

Tabelle 36 Color scheme IO density heat - meaning of the colors

Color	Status	Meaning
Red	Exceeded	Current average response time is above the threshold
Yellow	Vulnerable	Current average response time is 80% above the threshold
Green	Normal	Current average response time is less than 80% of the threshold
White	Filtered	RT due to insufficient IO (<10 IO / s) filtered by the scheduler
Faded	Undefined	Not enough values are known for the calculation : DRT, RT, or IOr = undef.
Object type	Not Applicable	Object (-type) is not affected

Status aggregation

- Higher grouping **storage consuming** object types, such as VDisk, VDisk group, data class ... aggregate the status based on the object type *VDisk copy*.
- Higher grouping **storage providing** object types, such as MDISK group, storage class ... aggregate the status based on the object type *MDISK.*

In these cases the "worst" status determines the aggregation result.

HowTo: usage of response time compliance coloring strategy

1. Setting the MDisk group attribute "style response time" (useful values: 1 to 30ms)
2. Selecting the following "predefined favorite": "**Analyse performance -> Response time compliance: Cluster ... MDg ... VDisk**"
3. Result: highlighting of VDisks which exceed a set response time in the treemap:

Abbildung 54 Color scheme design response time - screenshot of a compliance map

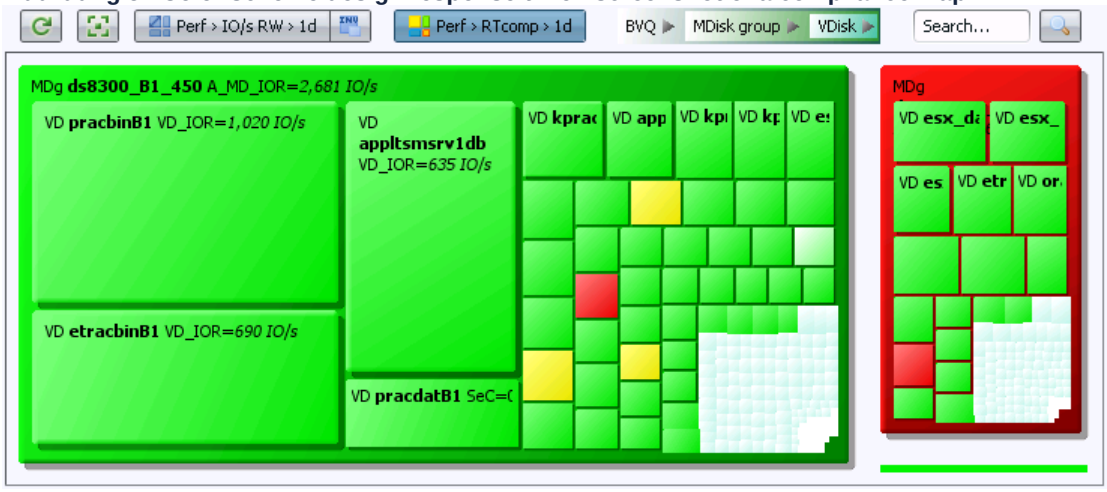

1.5.4.17 Color scheme overall compliance

The object staining according to the scheme "overall compliance" colorizes the objects of the treemap according to the combined status of all compliance conditions.

The compliance state "overall compliant" is calculated for every object type containing at least one of the data classes *mobility zone* or *VDG relations compliance state*s. Compliance conditions are combined and evaluated in this fashion.

In the treemap the compliance state is expressed by an adjusted color of the object area:

Tabelle 37 Color schema overall compliance - meaning of the colors

Color	Symbol	Condition	Meaning
Green		Compliant	Current configuration of the object meets all existing presets
Yellow		Incomplete	Specification of compliance targets (currently only VDisk group relation) is not complete
Red		Not compliant	Current configuration of the object does not meet all existing presets
Faded		Undefined	No specifications are defined for the object
Object type		Not applicable	Object (type) is not affected

Aggregation of c ompliance states

It should be noted, that the compliance state is determined only for a few objects directly . For most objects - especially higher grouping objects - the default state is "compliant" only if **all** relevant child objects have the status compliant.
For example, the target goal of an application's data class is only met, if all application´s VDisk groups meet their data class requirements.

Since all compliance requirements are optional, the compliance status "undefined" of an object is interpreted as "compliant" in the status aggregation of the container object. Exception: the aggregated state is "undefined", when all states are "undefined".

The state "not applicable" has no effect on the aggregation, unless all states are "not applicable", then it is also the aggregation's result.

Table: aggregation of compliance states

Condition object 1	Condition object 2	Condition aggregated

1.5.4.18 Color scheme data class compliance

The object color scheme "data class compliance" colors the objects of the treemap according to the SLA conformity, defined in the BVQ DC to SC placement rating.

The compliance state for the default data type class is determined for the object type VDisk copy. All other object types, which also display this condition type, determine its state by aggregation.

In the treemap this compliance state is expressed by following stainings of the object areas:

Tabelle 38 Color scheme data class compliance - meaning of the colors

Color	Symbol	Condition	Meaning
Green		Compliant	Requirements met: VDisk copy placement rating >=0 (or)
Yellow		Incomplete	Requirements definition not complete
Red		Not compliant	Requirements not met: VDisk copy Placement rating < 0 ( )
Faded		Undefined	No object-requirements defined
Object type		Not Applicable	Object (-type) is not affected

1.5.4.19 Color scheme mobility zone compliance

The object coloring according to the scheme "mobility zone compliance" colors the objects of the treemap on the compliance-state with requirements defined in the object type BVQ mobility zone.

This color scheme indicates whether VDisks currently meet the user requirements with regard to the restrictions of their movement clearance.

Examples:

- VDisks of the group "DB2 logs RZ1" may only lie on MDISK groups which in turn lie on subsystems in "RZ1".
- VDisks of the group "file system business critical" may only lie on MDISK groups which in turn lie on enterprise subsystems.

In the treemap this compliance state is expressed by the following coloring of the object areas :

Tabelle 39 Color scheme data class compliance - meaning of the colors

Color	Symbol	Condition	Meaning
Green		Compliant	The objective is met: the VDisk copies lie in the BVQ mobility zone, defined by the VDisk group
Yellow		Incomplete	Definition of the default(s) is incomplete
Red		Not compliant	The specification of the state "compliant" are not kept
Faded		Undefined	No specifications are defined for the object

Color	Symbol	Condition	Meaning
Object type		Not applicable	Object (-type) is not affected

1.5.4.20 Color scheme VDg relation compliance

The object coloring according to the scheme "VDG relation compliance" colors the objects of the treemap on the state of requirements´ compliance defined by the object type BVQ VDisk group relation.

This color scheme indicates, whether SVC VDisk actually complies with the requirements depending on other SVC VDisks.

Examples:

- SVC VDisks of the group "DB2 Logs RZ1" are only permitted to be placed in SVC MDisk groups, which in turn are placed on subsystems in "RZ1".
- SVC VDisk of the group "filesystem business critical" are only permitted to be placed in SVC MDisk groups, which in turn are placed on enterprise subsystems.

In the treemap this compliance state is expressed by the following surfaces´ colors:

Tabelle 40 Color scheme data class compliance - meaning of the colors

Color	Symbol	Condition	Meaning
Grün		Compliant	The objective is met: the VDisk copies are compliant with the conditions in the BVQ VDisk group relation.
Gelb		Incomplete	Specification is not complete
Rot		Not compliant	Specification of the condition "compliant" is not met
Faded		Undefined	No specifications are defined for the object
Object type		Not applicable	Object (type) is not affected

The compliance status for the default VDisk group relation type is determined for the object type VDisk copy (aggregation source), and separately determined for the VDisk group relation.
All other object types displaying the condition type determine its state by aggregation.

Symbol	Condition	Description
	Compliant	VDisk copy: the placement of the VDisk copies complies with the

Symbol	Condition	Description
		conditions of all assigned VDisk group relations
	Not compliant	VDisk copy: the placement of the VDisk copies violates the condition of one or more assigned VDisk group relations
	Incomplete	VDisk copy: at least one of the assigned VDisk group relations refers to undefined BVQ objects (room, site, ...)
	Compliant	VDisk group relation: the placement of all assigned VDisk copies complies with the requirement of the VDisk group relation
	Not compliant	VDisk group relation: the placement of one or more assigned VDisk copies violates the condition of the VDisk group relation
	Incomplete	VDisk group relation: at least one of the assigned VDisk copies lacks relations to BVQ objects that are required for the VDG rel (room, site, ...)

1.5.4.21 Object type specific color schemes

Some object types have a specially adapted color scheme for a better representation of their specific properties.

These are

- Color scheme SVC controller port, SVC node port, SVC host port
- Color scheme port connect
- Color scheme VM virtual machine

Color scheme port connect

The special coloring of the object type port connect colors the treemap´s objects depending on the connected port, and the active / inactice status of the connection.

Colors of connected SVC port types:

Tabelle 41 Meaning of the color object type port connect

Color	SVC node port is connected with
Green	SVC controller port
Turquoise	SVC node port

Color	SVC node port is connected with
Blue	SVC host port
Grey	SVC unkonwn port

The color appears darkened when the connection´s status is reported as "inactive" by the SVC (see also CLI command iscontroller).

Abbildung 55 Coloring example for object type

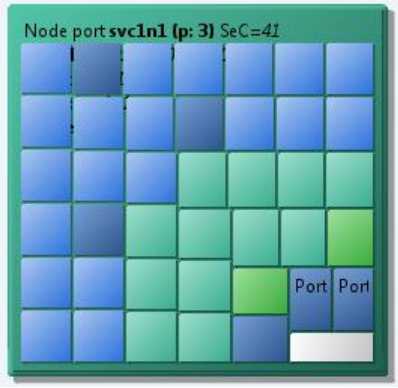

Please see also: color scheme SVC controller port, SVC node port, SVC host port.

Color scheme VM virtual machine

The special coloring of the object type VM virtual machine paints objects in the treemap, depending on their power status.

Color depending on power status:

Tabelle 42 Meaning of coloring object type

Color	Power status
Pale lilac	Power on
Faded lilac	Power off or suspended

Abbildung 56 Coloring example object type

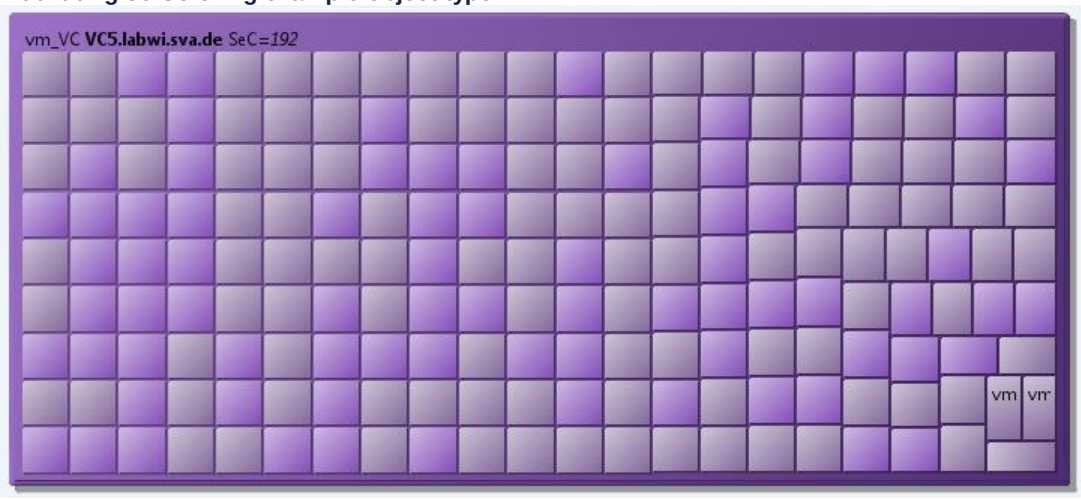

Color scheme SVC controller port, SVC node port, SVC host port

The special coloring of the object types SVC controller port, SVC node port and SVC host port paints the treemap´s objects depending on the port´s status.

Color per SVC port type and status:

Tabelle 43 Meaning of the object

Color	SVC port type / status
Green	SVC controller port
Turquoise	SVC node port and SVC unkonwn port
Blue	SVC host port
Yellow	Status "inactive"
Red	Status "offline"
Grey	Status "degraded"

The states "offline", "degraded" and "inactive" are derived from the output of the SVC CLI commands iscontroller, isnode, ishost and isfabric.

Abbildung 57 Coloring example for the object type

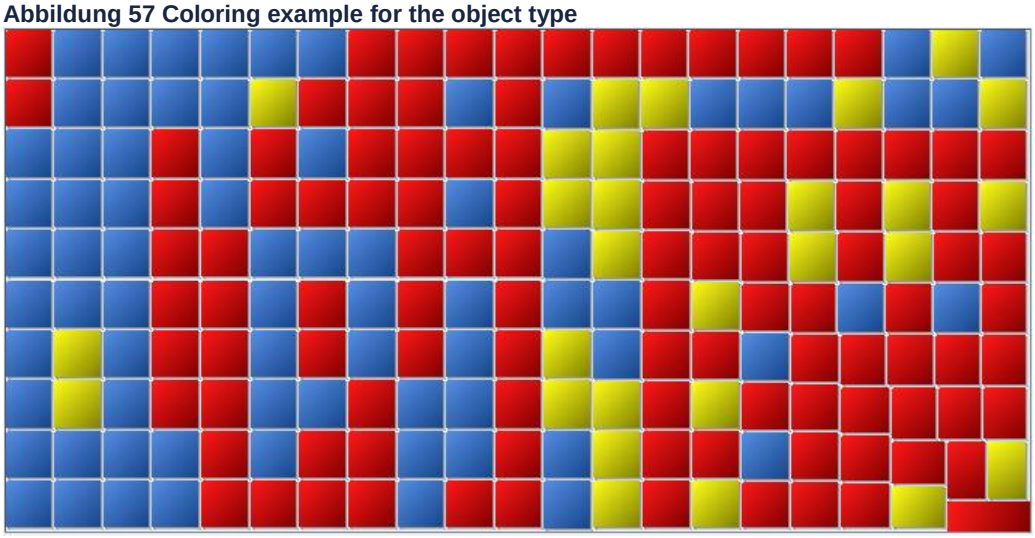

See also: color scheme port connect

1.5.5 Topology - search function

Via the search function, objects can be found and specifically selected in the treemap.

The search field of the topology-view is also accessible via the following shortcut:

[Ctrl f]	Opens the treemap´s search field

In the search Field you can e.g. enter fractions of the name of displayed objects, but also all other object-attributes are scanned.

The entry is confirmed through „Enter" or the clicking of the Loupe-Symbol.

1.5.5.1 Hit list- search dialog

If there are multiple hits, the *search dialog* box opens. It contains all hits (objects) and the appropriate attribute.

Abbildung 58 Search field result list:

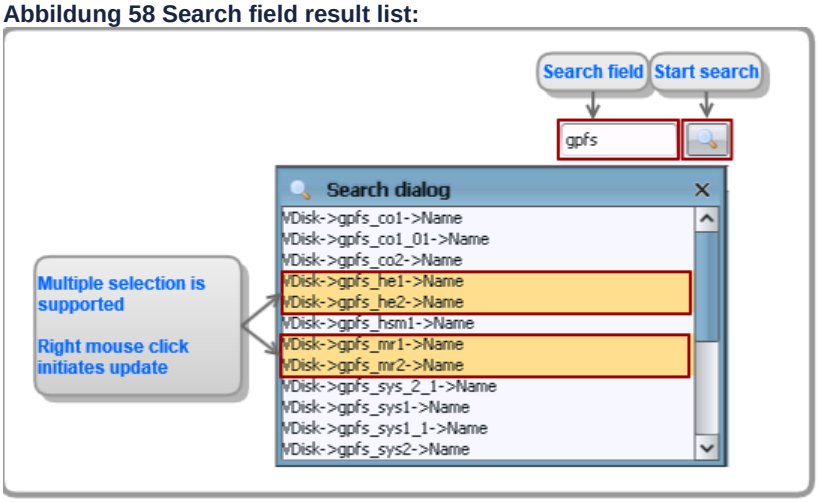

Selection, zoom and highlighting of correlated objects in the treemap.

The objects on the hit list are automatically marked as "selected" in the treemap. The treemap in addition is enlarged (zoom), to display all found objects as large as possible.

A selection of objects in the hit list leads to additional highlighting (flashing) of these objects in the treemap, to find their positions more easily.

By using the usual multi-select keys [Ctrl], [Shift] or [Ctrl + A] more than one object can be selected in the results list.
A final double-click and holding down the multi-select button at the same time leads to the selection and zoom of correlating objects in the treemap.

With a single click on the right mouse button, all selected objects in the treemap are prompted to light up again.
By double clicking the right mouse button, the treemap is zoomed on all selected objects.

With [Esc] the hit list can be closed again.

Operation´s summary

Action / button	Selection in result list	Highlighting in treemap	Selection and zoom in treemap
Single left click	✓	✓	✗
Double left click	✓	✓	✓
Single right click	✗	✓	✗
Double right click	✗	✓	✓
[Ctrl], [Shift] + left click	Multiple objects	✓	✗

Action / button	Selection in result list	Highlighting in treemap	Selection and zoom in treemap
[Ctrl+A]	All objects	✗	✗
[Esc]	Closes the hit list	--	--

1.5.5.2 Search with wildcards

The search function can be expanded by the usage of wildcards.
The following placeholders are available as wildcards:

?	Question mark	Exactly one character
*	Star key	Any number of characters

1.6 Performance view

In the performance view load-curves for selected objects can be displayed in a chart.

Abbildung 59 Overview performance-view

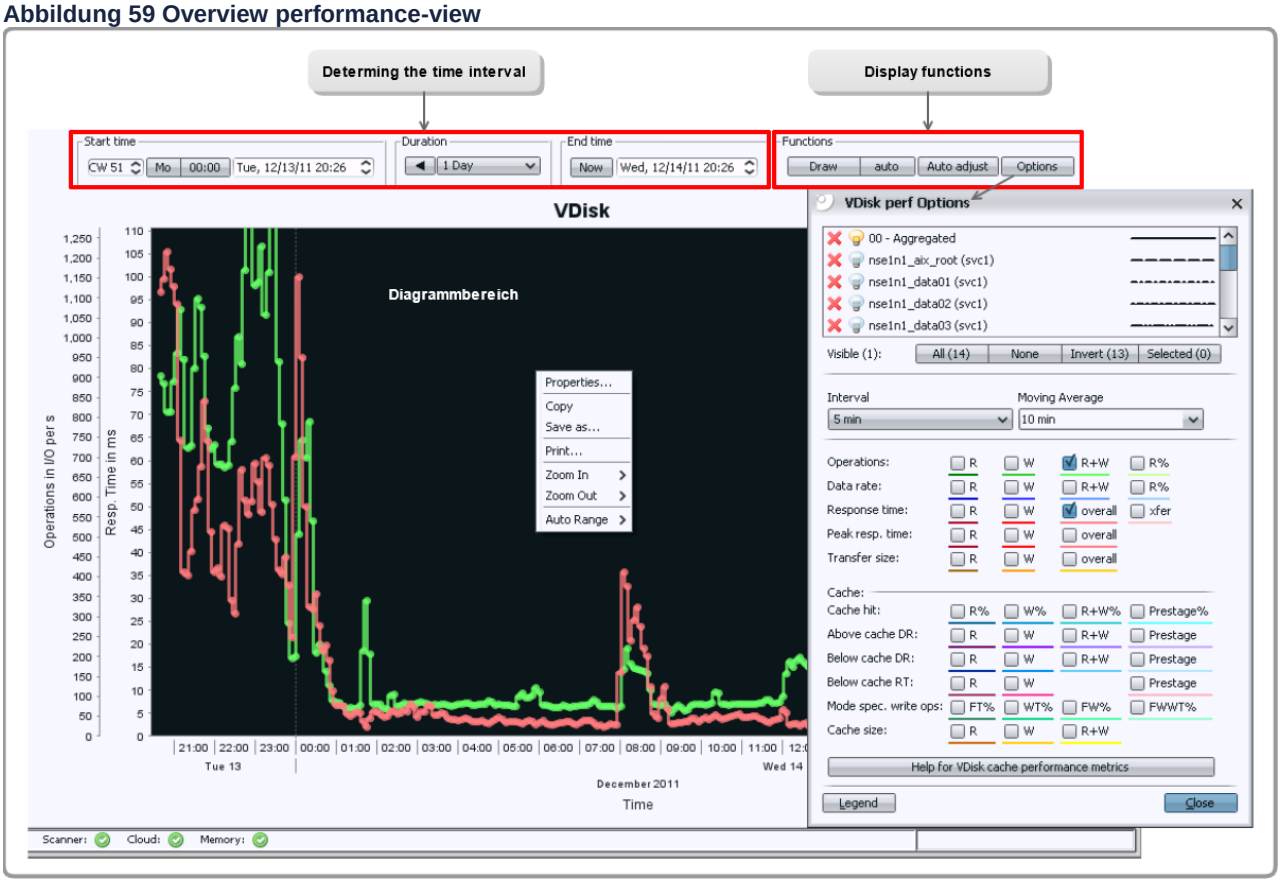

1.6.1 Main panels

Tabelle 44 Performance-view: main panels

Purpose	Panels	Functions
Selection of the time period	Start time, duration, end time	Allocation of the visible time period the chart area of the performance view by choosing start time, duration, and end time.
Loading and Drawing	Draw, manual/auto	Automatically or manually reloading and drawing of the the load curves in the diagram area of the performance view.
Scaling of the value axes	Auto adjust	Manual triggering of the value axes´ automatic scaling in the graph area of the performance view.
Objects and metrics	Options dialog	The dialog *options* in the performance view allows you to select the objects and their metrics to be displayed as a load curve in the chart area.
Display load	Chart area	In the chart area of the performance view the load curves

Purpose	Panels	Functions
curves		of selected objects and performance metrics are displayed.

1.6.2 Walkthrough

1.	Calling	Select objects in the treemap and call the performance view via the pop-up menu
2.	Selecting time period	Set the desired time period
3.	Selecting metrics	Select the desired metrics and objects
4.	Analyzing chart	Now the graph can be interpreted, zoomed, printed, exported ...

1.6.3 Call of performance view

Performance view is usually called from the object context menu of the treemap objects. If multiple objects are selected, they are displayed in a shared performance view. It is recommend not to select too many objects at once, otherwise the loading and drawing speed slows down and the GUI´s memory consumption increases.

Abbildung 60 Call of performance view

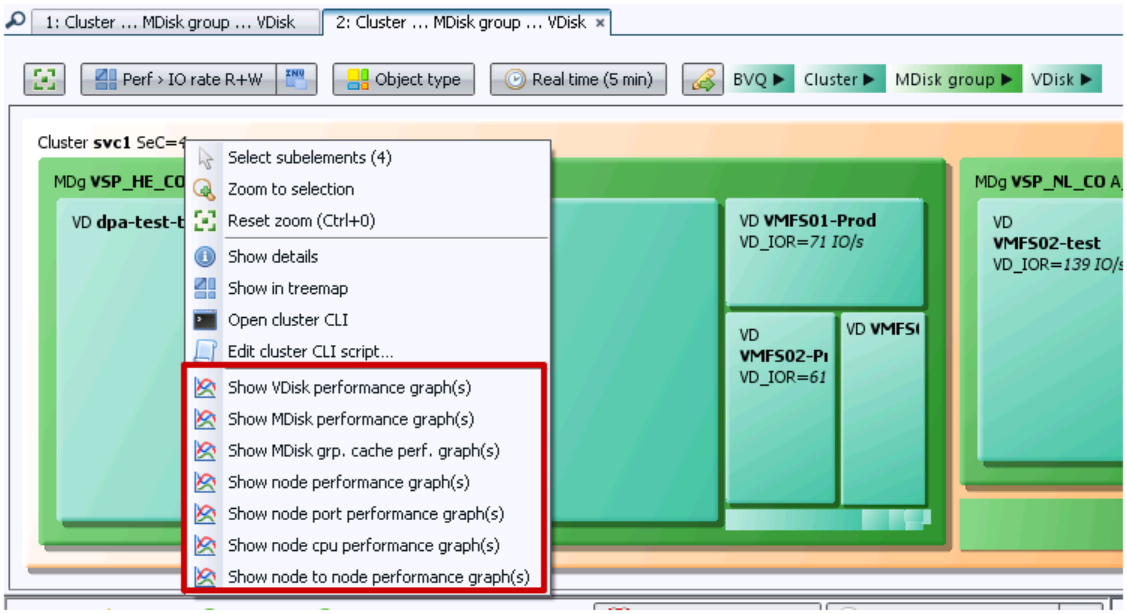

1.6.3.1 Basic objects und grouping objects

From the storage system load values for the following four basic object types are transferred to BVQ:
VDisk (volume), MDisk, node, node port

Almost all BVQ objects arrange linked objects of one or multiple basic types.
For these objects, a useful performance view can also be derived via the aggregation (sum or average generation) of the associated objects.

Therefore, the following options exist to call the performance view for each object type:

Tabelle 45 Basic objects und grouping objects

Object type	VDisk	MDisk	Node	Node port
VDisk group	✓			
VDisk	✓	✓	✓	✓
VDisk copy	✓	✓		
MDisk group	✓	✓		
MDisk	✓	✓		
Cluster	✓	✓	✓	✓
IO group			✓	✓
Node	✓		✓	✓
Node port			✓	✓
Port connect			✓	
Host	✓		✓	
Host port			✓	
Controller (SVC)	✓	✓		
Controller	✓	✓		
Controller port		✓	✓	
Disk Array	✓	✓		
Disk type		✓		

Object type	VDisk	MDisk	Node	Node port
Application	✓			
Application group	✓			
Cost center	✓			
SLA	✓			
Data class	✓			
Mobility zone	✓	✓		
Room	✓	✓	✓	
Site	✓	✓	✓	

1.6.4 Selection of the time period

Allocation of the visible time period the chart area of the performance view by choosing start time, duration, and end time.

Abbildung 61 Performance view: selection of the time period

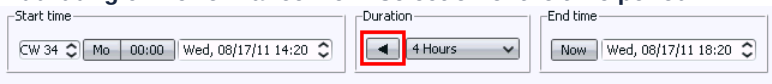

Tabelle 46 Performance view: functions of the time period´s selection

Categorie	Element	Functions´ description
Start time	CW	Calendar week - allows to jump forwards or backwards by the week
Start time	Mo	Monday - sets the start time to the Monday of the selected
Start time	00:00	00:00 - Sets the start time to the beginning of the selected day
Start time	Date	Allows the entry of the start time *
Duration	Duration	Selecting the width (duration) of the displayed interval
Duration	< >	Effective direction of the set duration
End time	Now	Now - sets the end time to the current time
End time	Date	Allows the entry of the end time *

* In the fields "date" a double click on the Hour digits denoting the hour select the hours of time. A subsequent click e.g. on "^" increases the time by one hour. This also works with the digits denoting month, day, year and minutes.

The displayed time period is always determined by the following three controls:

Start time	Duration	End time

If one of the three is changed, the other two will be adjusted automatically.
This allows to move the time period with a single click on the timeline.

1.6.4.1 Effective direction of the selected "duration "

The current effective direction is indicated by an arrow button to the right or left of the selected "duration".
By clicking on the arrow the effective direction can be reversed at any time.
As a rule, it now points from the end time to start time. An increase or decrease keeps the right end point (no need to click on "new").
The effective direction always points **away** from the last changed field (start/end) to the other.

1.6.5 Loading and drawing of the graph

Automatically or manually reloading and drawing of the the load curves in the diagram area of the performance view.

Abbildung 62 Performance view: functions for loading and drawing of the load curves

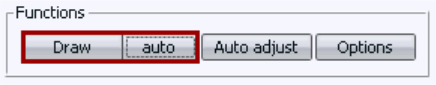

Draw	Triggers manually a new loading/ drawing
Auto / manual	Sets manually or automatically new loading/drawing

Normally, the loading and drawing of individual load curves is sufficiently fast. Therefore, this is done automatically after each change of the time period.
If this is too slow, or if multiple entries are required to select the time period, this automatic can be canceled by using the "auto / manual".

1.6.6 Scaling of the value axes

Manual triggering of the value axes´ automatic scaling in the graph area of the performance view.

The scaling of the value axes is deliberately kept constant when the selected time period is changed, to prevent unwanted "discontinuities" of the load curves by choosing different time sections.
The button "auto adjust" starts an optimal re-scaling of the value axes.

Abbildung 63 Performance view: auto-scaling of the value axes - auto adjust

Auto adjust	Sets the maximum value of all values axis (y-axis) to the maximum value of the displayed load curve

1.6.7 Options dialog for objects and metrics

The dialog *options* in the performance view allows you to select the objects and their metrics to be displayed as a load curve in the chart area.

1.6.7.1 Calling the options dialog

To configure the displayed objects and their metrics the button **options** of the options dialog can be used.

Abbildung 64 Performance view: button to call the options dialog

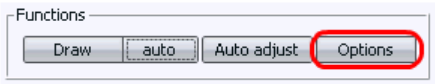

1.6.7.2 Main areas of the options dialog

The options dialog has the following main areas:

Abbildung 65 Performance view: dialog

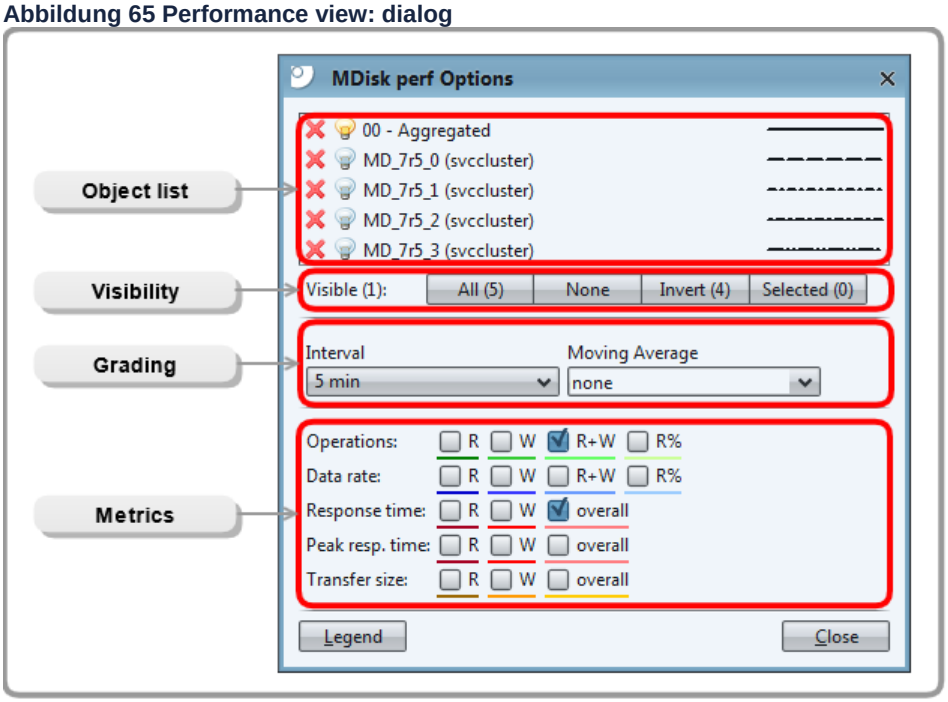

Tabelle 47 Main areas of the options dialog

Object list	Contains selected objects from this performance view and optionally the aggregation object
Visibility	Affects the visibility of multiple objects in the chart´s object list
Grading	Enables a smoothing of the load curves with *interval widening* and *moving average*
Metrics	Selection of object type-specific metrics to display them in the chart
Other	Turning on/off the legend below the chart; if necessary calling of help pages

Object list

The object list maintains all loaded objects of the performance view in alphanumerically sorted order.

A special stroke type is assigned to each object, which is repeated every 5 objects.

The following functions are available in this list:

Tabelle 48 Object list - functions and notes

✖ (Remove)	Removes objects from the context (aggregation object is created automatically)

(Fade-in, Fade-out)	Fading in and out in the graph area
Line type	Line type of the object in the chart
Tooltip	A short pause of the cursor opens the tooltip with short-term performance indicators
Multi select	By pressing the [Shift] or [Ctrl] key multiple objects can be selected
Selection bold	Selected objects are shown in bold for better visibility in the graph
Context menu functions	• Show details – calls the object detail view of the selected object • Export – exports curve values of the object in CSV format (e.g. for import into MS Excel) • Edit – calls to the appropriate object editor for this object type.

Aggregated load curve

If more than one object is included in the list the aggregation-object **00-aggregated** is displayed additionallly at the top.

Abbildung 66 MDisk performance options - selection of an aggregated load curve

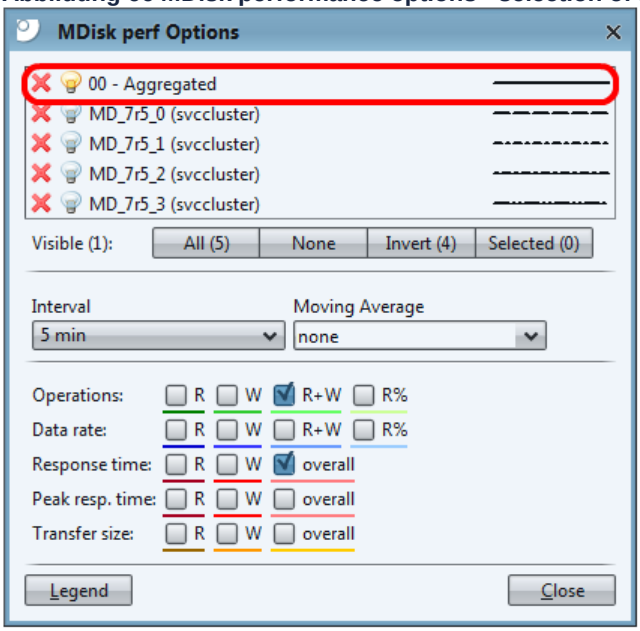

In the graph area, the aggregation-object represents the sum or average value curves for all other objects in the list.
It has the same features as the other objects in the list.

Application: chart the overall course of a MDISK group´s or controller´s IO rate.

Visibility

The following four switches affect the visibility of multiple objects simultaneously:

Tabelle 49 Visualization-functions for visibility of objects

Visible (n)	Number of currently visible objects
All (n_a)	Makes all (n_a) objects visible
None	Makes all objects invisible
Invert (n_i)	Reverses the current visibility of all objects. Then (n_i) objects are visible.
Selected (n_s)	Makes all currently selected (n_s) objects visible

Application: Quickly change between aggregation curve and all object curves via switches [invert].

Smoothing procedures

To smooth the displayed curves, 2 methods are available:

Tabelle 50 Smoothing procedures

Interval	Interval extension - uses a arithmetic mean to combine multiple data points
Moving average	Smoothing method for the left-sided moving average

Both methods can be usefully combined.

- The smoothing by using interval extension leads to a reduction of the curve nodes - well suited for a very wide period of time with many points
- The smoothing by using moving average, keeps the number of curve nodes - suitable to highlight a course better by trend

Metrics

The metrics of the performance view are computed from one or more measured metrics (performance statistics) of an underlying basic object type.

To each measurement metric belongs a time series of measurements collected and delivered by the storage system (SVC, V7000, ...).

Directly below the smoothing procedure panels, the selection area for metrics can be found. Any combination of metrics can be selected by using check boxes.

Each metric has its distinct color, which is shown in the line below the box.
A contiguous group of metrics has the same basic color, but with different levels of brightness.

Depending on the object type, the performance metrics are further divided into chapters.

To each box, a tooltip can be displayed, which contains more information about each metric. It also specifies the formula (strongly simplified) of the underlying measurement metrics.

In the chapter performance metrics a detailed description about metrics can be found.

Other points

Legend	Toggles the legend below the graph
Help	Help for complex metrics (currently only VDisk cache)

The legend is displayed in the graph area and can take a lot of space there when many objects and metrics are selected. Therefore, it is disabled by default.

1.6.8 Chart area

In the chart area of the performance view the load curves of selected objects and performance metrics are displayed.

1.6.8.1 Representation of the load curves

The objects and metrics to be displayed, are selected in the options dialog of the performance view.

Each object has a different line type, each metric a different color.

Multiple metrics with different unit of measurement on different axes can be displayed in the same chart.

A grey area displays possible missing values (holes) in the chosen time period.

The displayed time section and the reloading behaviour is selected with the panels above the chart area.

Below the time axis a legend can be faded in via the options dialog.

1.6.8.2 Zooming of the curves in the graph area

Zoom in/out	Selection of an area on the graph increases it

Zoom in: by drawing a selection rectangle from top left to bottom right over a chart area, it is enlarged.

Zoom out: by drawing a selection rectangle from bottom right to top left corner over an area of load curves, the diagram is restored to normal scaling.

Please see also: scaling of the value axes

1.6.8.3 Pop-up menu of the chart area

The popup menu is opened by right-clicking in the graph area and allows the selection of the following functions:

Tabelle 51 Performance view: context menu functions of the chart area

Properties	Opens chart properties dialog box to edit the chart settings
Copy	Copies the chart area with axes-labeling
Save as	Saves chart area with axes-labeling as a PNG-image.
Print	Prints the chart area with axes-labeling
Zoom in	Zoom in function, selection: both axes, domain axes or range axes
Zoom out	Zoom out function, selection: both axes, domain axes or range axes
Auto scale	Scales to output size, selection: both axes, domain axes or range axes

1.6.8.4 Tooltip on measuring points

A short pause of the mouse pointer on a measuring point opens a tooltip which displays the currently measured values for all metrics of this object type.

Abbildung 67 Performance view: tooltip of the chart area

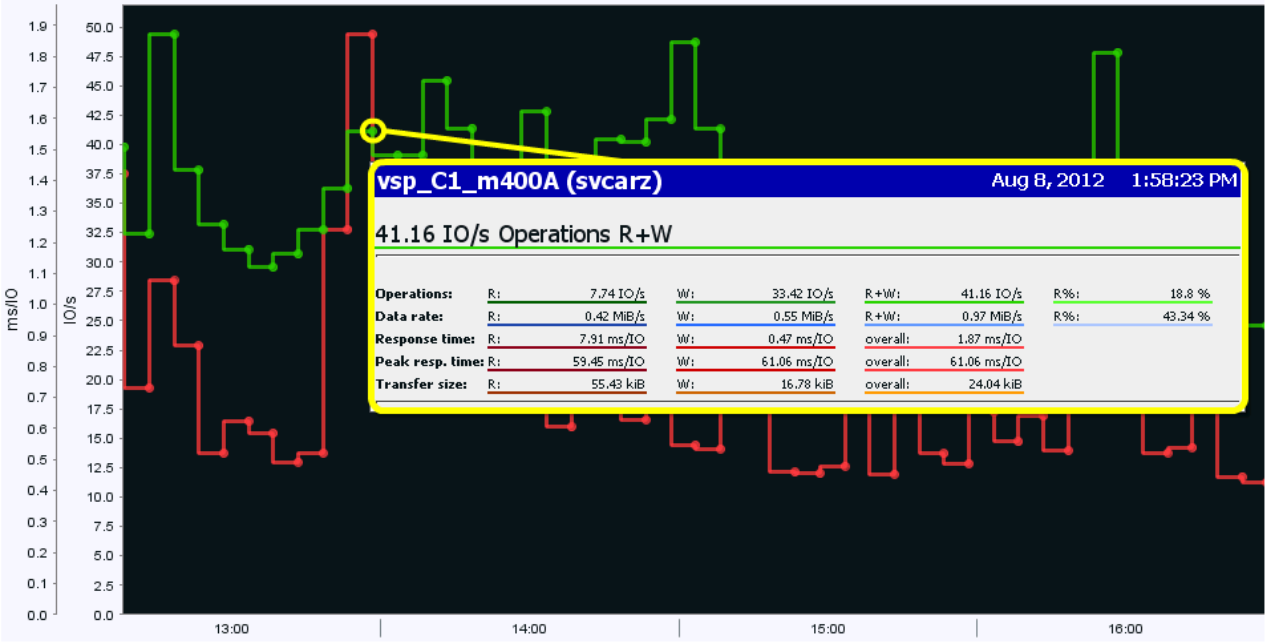
vsp_C1_m400A (svcarz)
Aug 8, 2012 1:58:23 PM
41.16 IO/s Operations R+W
Operations: R: 7.74 IO/s W: 33.42 IO/s R+W: 41.16 IO/s R%: 18.8 %
Data rate: R: 0.42 MiB/s W: 0.55 MiB/s R+W: 0.97 MiB/s R%: 43.34 %
Response time: R: 7.91 ms/IO W: 0.47 ms/IO overall: 1.87 ms/IO
Peak resp. time: R: 59.45 ms/IO W: 61.06 ms/IO overall: 61.06 ms/IO
Transfer size: R: 55.43 kiB W: 16.78 kiB overall: 24.04 kiB
ms/IO
IO/s
13:00
14:00
15:00
16:00

1.6.9 Performance metrics

The metrics of the performance view are computed from one or more measured metrics (performance statistics) of an underlying basic object type.
To each measurement metric belongs a time series of measurements collected and delivered by the storage system (SVC, V7000, ...).

The selection of metrics to represent a load curve in the chart area of the performance view is carried out in the options dialog for objects and metrics.

The metrics are calculated by using a mathematical function f whitch has 1 .. n measurement metrics as variables:

Metric = f (measuremend $metric_1$, ... measurement $metric_n$)

For the description of the underlying basic measured metrics of the object types, please see the IBM documentation.

Measured metrics are captured by SVC and BVQ for the following object types:

- SVC VDisk
- SVC MDisk
- SVC node

In addition the following (abstract) performance object types are derived from the object type SVC node:

- MDisk group cache partition
- Node port
- Node to node

The following table contains a list of all available performance object types:

Tabelle 52 Overview performance object types

Performance Object type	Description
VDisk	Metrics of the object type *VDisk* which can be displayed as charts in the perfomance view.
MDisk	Metrics of the object type *SVC MDisk*, which can be represented in the perfomance view as load curve.
Node	Metrics of the object type *node*, which can be displayed as load curves in the perfomance view.
Node port	Metrics of the performance-object type *node port*, derived from the object type *SVC node*.
MDisk group cache partition	Metrics of the performance-object type *MDisk group cache partition* derived from the object type *SVC node*.
Node to node	Metrics of the performance-object type *node to node*, are derived from the object type *SVC node*.

1.6.9.1 VDisk performance metrics

Metrics of the object type *VDisk* which can be displayed as charts in the perfomance view.

Abbildung 68 Screenshot: options dialog: VDisk perf options

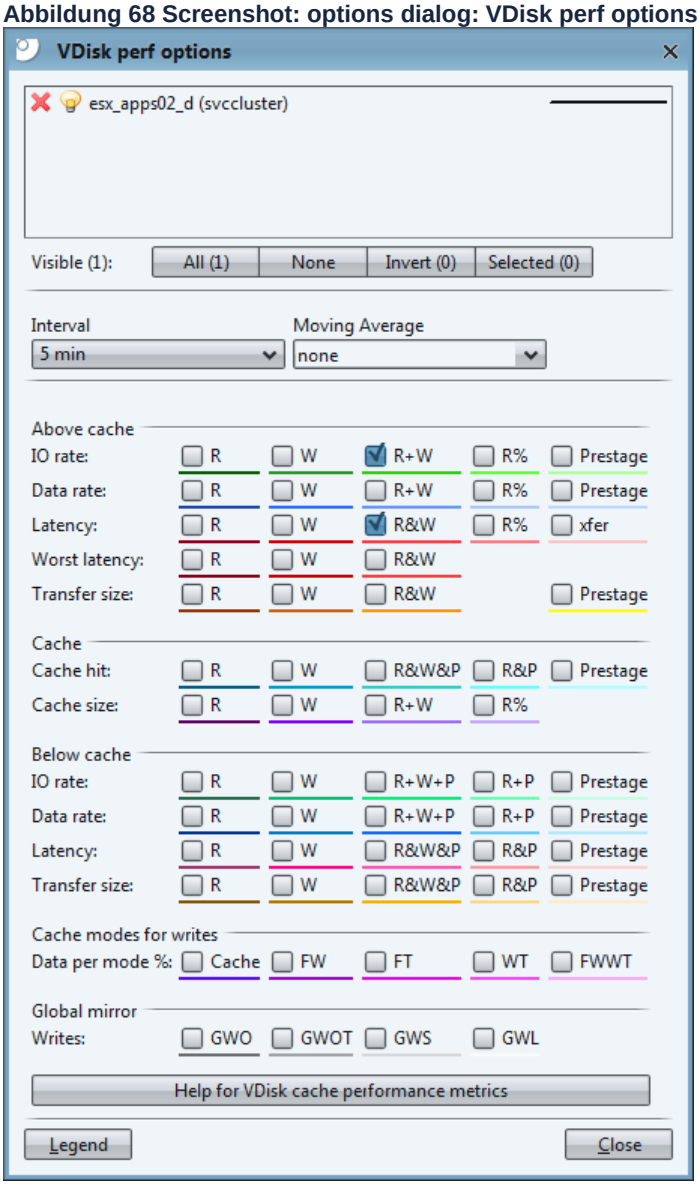

Tabelle 53 Overview: performance metrics VDisk

Above cache					
IO rate:	✅ R	✅ W	✅ R+W	✅ R%	✅ Prestage
Data rate:	✅ R	✅ W	✅ R+W	✅ R%	✅ Prestage
Latency:	✅ R	✅ W	✅ R&W	✅ R%	✅ xfer

Above cache					
Worst latency:	R	W	R&W		
Transfer size:	R	W	R&W		Prestage
Cache					
Cache hit %:	R	W	R&W&P	R&P	Prestage
Cache size:	R	W	R+W	R%	
Below cache					
IO rate:	R	W	R+W+P	R+P	Prestage
Data rate:	R	W	R+W+P	R+P	Prestage
Latency:	R	W	R&W&P	R&P	Prestage
Transfer size:	R	W	R&W&P	R&P	Prestage
Cache modes for writes					
Data per mode %:	Cache	FW	FT	WT	FWWT
Global mirror					
Writes:	GWO	GWOT	GWS	GWL	

Tabelle 54 Description: performance metrics VDisk

Metric	Unit Axis	Description
Above cache		
IO rate: R	ops/s	RO: read operation rate (ops/s)
IO rate: W	ops/s	WO: write operation rate (ops/s)
IO rate: R+W	ops/s	RO+WO: total operation rate (ops/s)
IO rate: R%	%	RO/(RO+W100O): % reads of total operations (%)
IO rate:	ops/s	CTRHP: read operation rate on prestaged data (ops/s)

Metric	Unit Axis	Description
Prestage		
Data rate: R	MiB/s	RB *512: read data rate (MiB/s)
Data rate: W	MiB/s	WB *512: write data rate (MiB/s)
Data rate: R+W	MiB/s	RB+WB *512: total data rate (MiB/s)
Data rate: R%	%	RB/(RB+WB): % read data of total data (%)
Data rate: Prestage	MiB/s	CTRHPs *512B: read data rate on prestaged data (MiB/s)
Latency: R	ms/op	RL/RO: average response time per read operation (ms/op)
Latency: W	ms/op	WL/WO: average response time per write operation (ms/op)
Latency: R&W	ms/op	(RL+WL)/(RO+WO): average response time per operation (ms/op)
Latency: R%	%	RL/(RL+WL): average % read response time of total response time (%)
Latency: xfer	ms/op	XL/(RO+WO): average transfer latency per operation (ms/op)
Worst latency: R	ms/op	RLW: response time of the worst read operation (ms/op)
Worst latency: W	ms/op	WLW: response time of the worst write operation (ms/op)
Worst latency: R&W	ms/op	MAX(RLW,WLW): response time of the worst operation (ms/op)
Transfer size: R	KiB/op	RB/RO: average transfer size per read operation (KiB/op)
Transfer size: W	KiB/op	WB/WO: average transfer size per write operation (KiB/op)
Transfer size: R&W	KiB/op	(RB+WB)/(RO+WO): average transfer size per operation (KiB/op)
Transfer size: Prestage	KiB/op	CTRHPs*512B / CTRHP: average transfer size per read operation on prestaged data (KiB/op)
Cache		

Metric	Unit Axis	Description
Cache hit %: R	%	CTRHs/CTRs: % cache hits on non prestaged data of read data (%)
Cache hit %: W	%	CTWHs/CTWs: % cache hits on dirty data of write data (re-writes before destage) (%)
Cache hit %: R&W&P	%	(CTRHs+CTWHs+CTRHPs)/(CTRs+CTWs+CTPs): % cache hits on all valid data of read and write data (%)
Cache hit %: R&P	%	(CTRHs+CTRHPs)/(CTRs+CTPs): % cache hits on all valid data of read data (%)
Cache hit %: Prestage	%	CTRHPs/CTRs: % cache hits on prestaged (read ahead) data of read data (%)
Cache size: R	MiB	CV-CM *512B: clean cache data (MiB)
Cache size: W	MiB	CM *512B: modified (dirty) write cache data (MiB)
Cache size: R&W	MiB	CV *512B: valid - read, prestaged & write - cache data (MiB)
Below cache		
IO rate: R	ops/s	CTRM: read cache miss operation rate (ops/s)
IO rate: W	ops/s	CTD: write cache destage operation rate (ops/s)
IO rate: R+W+P	ops/s	CTRM+CTD+CTP: total cache operation rate (ops/s)
IO rate: R+P	ops/s	CTRM+CTP: total read cache operation rate (ops/s)
IO rate: Prestage	ops/s	CTP: read ahead cache prestage operation rate (ops/s)
Data rate: R	MiB/s	CTRMs *512B: read cache miss data rate (MiB/s)
Data rate: W	MiB/s	CTDs *512B: write cache destage data rate (MiB/s)
Data rate: R+W+P	MiB/s	(CTRMs+CTDs+CTPs) *512B: total cache data rate (MiB/s)
Data rate: R+P	MiB/s	(CTRMs+CTPs) *512: total read cache data rate (MiB/s)
Data rate: Prestage	MiB/s	CTPs *512B: read ahead cache prestage data rate (MiB/s)

Metric	Unit Axis	Description
Latency: R	ms/op	SLAV/CTRM: average response time per read cache miss (demand stage) operation (ms/op)
Latency: W	ms/op	DLAV/CTD: average response time per write cache destage operation (ms/op)
Latency: R&W&P	ms/op	(SLAV+DLAV+PLAV)/(CTRM+CTD+CTP): average response time per total cache operation (ms/op)
Latency: R&P	ms/op	(SLAV+PLAV)/(CTRM+CTP): average response time per total read cache operation (ms/op)
Latency: Prestage	ms/op	PLAV: average response time per read ahead cache prestage operation (ms/op)
Transfer size: R	KiB/op	CTRMs/CTRM: average transfer size per read cache miss (demand stage) operation (KiB/op)
Transfer size: W	KiB/op	CTDs/CTD: average transfer size per write cache destage operation (KiB/op)
Transfer size: R&W&P	KiB/op	(CTWMs+CTDs+CTPs)/(CTWM+CTD+CTP): average transfer size per total cache operation (KiB/op)
Transfer size: R&P	KiB/op	(CTWMs+CTPs)/(CTWM+CTP): average transfer size per total read cache operation (KiB/op)
Transfer size: Prestage	KiB/op	CTPs/CTP: average transfer size per read ahead cache prestage operation (KiB/op)
Cache modes for writes		
Data per mode %: Cache	%	CTWs/WB: % of write data entering the cache layer (%)
Data per mode %: FW	%	CTWFWs/CTWs: % of write data handled in fast write cache mode (%)
Data per mode %: FT	%	CTWFTs/CTWs: % of write data handled in flush trough cache mode (%)
Data per mode %: WT	%	CTWWTs/CTWs: % of write data handled in write trough mode (%)

Metric	Unit Axis	Description
Data per mode %: FWWT	%	CTWFWSHs/CTWs: % of write data handled in write trough mode due to lack of memory (%)
Global mirror		
Writes: GWO	ops/s	GWO: Global mirror: overlapping write rate (w_ops/s)
Writes: GWOT	ops/s	GWOT: Global mirror: total (fixed or unfixed) overlapping write rate (w_ops/s)
Writes: GWS	ops/s	GWS: Global mirror: writes issued to the secondary site rate (w_ops/s)
Writes: GWL	ms/op	GWL/GWS: Global mirror: writes issued to the secondary site latency (ms/w_op)

1.6.9.2 MDisk performance metrics

Metrics of the object type *SVC MDisk*, which can be represented in the perfomance view as load curve.

Abbildung 69 Screenshot: dialog MDisk perf options

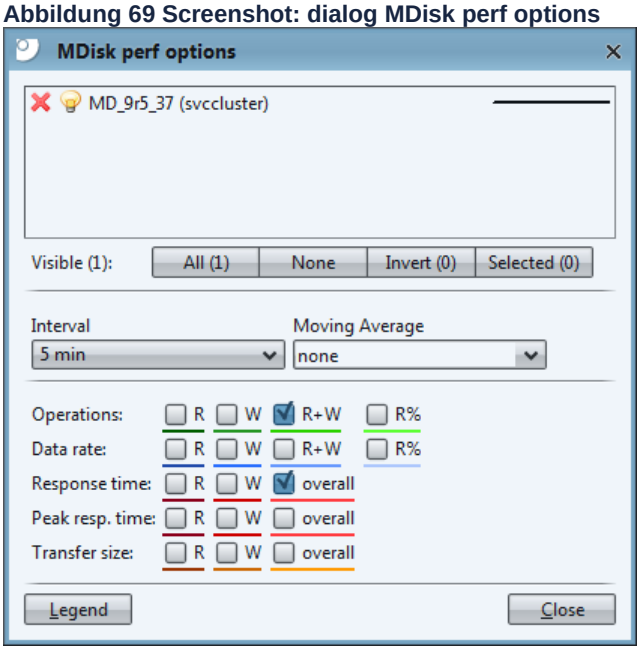

Tabelle 55 Overview: performance metrics MDisk

Metrics				
IO rate:	✅ R	✅ W	✅ R+W	✅ R%
Data rate:	✅ R	✅ W	✅ R+W	✅ R%
Latency:	✅ R	✅ W	✅ R&W	
Worst latency:	✅ R	✅ W	✅ R&W	
Transfer size:	✅ R	✅ W	✅ R&W	

Tabelle 56 Description: performance metrics MDisk

Metric	Unit Axis	Description
IO rate: R	ops/s	RO: Read operation rate (ops/s)
IO rate: W	ops/s	WO: Write operation rate (ops/s)
IO rate: R+W	ops/s	RO+WO: Total operation rate (ops/s)
IO rate: R%	%	RO/(RO+W100O): % Reads of total operations (%)
IO rate: prestage	ops/s	CTRHP: Read operation rate on prestaged data (ops/s)
Data rate: R	MiB/s	RB *512: Read data rate (MiB/s)
Data rate: W	MiB/s	WB *512: Write data rate (MiB/s)
Data rate: R+W	MiB/s	RB+WB *512: Total data rate (MiB/s)
Data rate: R%	%	RB/(RB+WB): % Read data of total data (%)
Data rate: prestage	MiB/s	CTRHPs *512B: Read data rate on prestaged data (MiB/s)
Latency: R	ms/op	RL/RO: Average response time per read operation (ms/op)
Latency: W	ms/op	WL/WO: Average response time per write operation (ms/op)
Latency: R&W	ms/op	(RL+WL)/(RO+WO): Average response time per operation (ms/op)
Worst latency: R	ms/op	RLW: Response time of the worst read operation (ms/op)

Metric	Unit Axis	Description
Worst latency: W	ms/op	WLW: Response time of the worst write operation (ms/op)
Worst latency: R&W	ms/op	MAX(RLW,WLW): Response time of the worst operation (ms/op)
Transfer size: R	KiB/op	RB/RO: Average transfer size per read operation (KiB/op)
Transfer size: W	KiB/op	WB/WO: Average transfer size per write operation (KiB/op)
Transfer size: R&W	KiB/op	(RB+WB)/(RO+WO): Average transfer size per operation (KiB/op)
Transfer size: prestage	KiB/op	CTRHPs*512B / CTRHP: Average transfer size per read operation on prestaged data (KiB/op)

1.6.9.3 Node performance metrics

Metrics of the object type *node*, which can be displayed as load curves in the perfomance view.

Abbildung 70 Screenshot: dialog node perf options

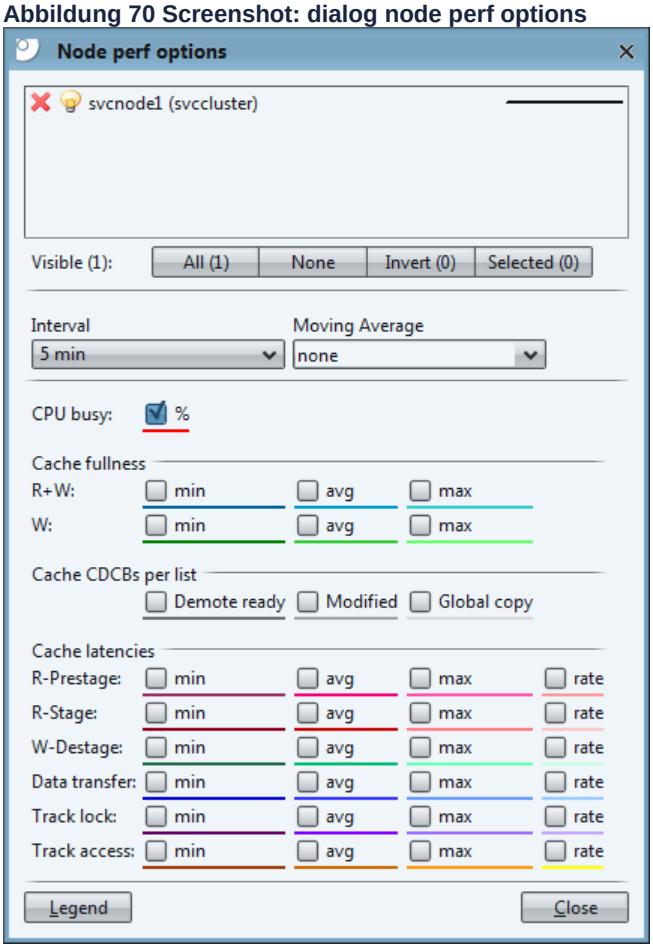

Tabelle 57 Overview: performance metrics node

CPU				
CPU busy	✔ %			
Cache fullness				
R+W:	✔ min	✔ avg	✔ max	
W:	✔ min	✔ avg	✔ max	
Cache CDCBs per list				
	✔ Demote ready	✔ Modified	✔ Global copy	
Cache latencies				

CPU				
R-Prestage:	min	avg	max	rate
R-Stage:	min	avg	max	rate
W-Destage:	min	avg	max	rate
Data transfer:	min	avg	max	rate
Track lock:	min	avg	max	rate
Track access:	min	avg	max	rate

Tabelle 58 Description: performance metrics node

Metric	Unit Axis	Description
Cache fullness		
R+W min	%	CFMN: Read and write cache fullness minimum (%)
R+W avg	%	CFAV: Read and write cache fullness average (%)
R+W max	%	CFMX: Read and write cache fullness maximum (%)
W min	%	WCFMN: Write cache fullness minimum (%)
W avg	%	WCFAV: Write cache fullness average (%)
W max	%	WCFMX: Write cache fullness maximum (%)
Cache CDCBs per list		
Demote ready	count	DRLL: Number of CDCBs for unmodified data tracks in the demote ready list (count)
Modified	count	MLL: Number of CDCBs for modified data tracks to be destaged in the modified list (count)
Global copy	count	GCLL: Number of CDCBs for data tracks in the global copy list (count)
Cache latencies		

Metric	Unit Axis	Description
W-Destage min	ms/op	DLMN: Destage latency minimum (ms/op)
W-Destage avg	ms/op	DLAV: Destage latency average (ms/op)
W-Destage max	ms/op	DLMX: Destage latency maximum (ms/op)
W-Destage rate	ops/s	DLCN: Destage operation rate (ops/s)
R-Prestage min	ms/op	PLMN: Prestage (read ahead) latency minimum (ms/op)
R-Prestage avg	ms/op	PLAV: Prestage (read ahead) latency average (ms/op)
R_Prestage max	ms/op	PLMX: Prestage (read ahead) latency maximum (ms/op)
R-Prestage rate	ops/s	PLCN: Prestage (read ahead) operation rate (ops/s)
R-Stage min	ms/op	SLMN: Stage (read miss) latency minimum (ms/op)
R-Stage avg	ms/op	SLAV: Stage (read miss) latency average (ms/op)
R-Stage max	ms/op	SLMX: Stage (read miss) latency maximum (ms/op)
R-Stage count	ops/s	SLCN: Stage (read miss) operation rate (ops/s)
Data transfer min	ms/op	DTMN: Data transfer latency minimum (ms/op)
Data transfer avg	ms/op	DTAV: Data transfer latency average (ms/op)
Data transfer max	ms/op	DTMX: Data transfer latency maximum (ms/op)
Data transfer rate	ops/s	DTCN: Data transfer operation rate (ops/s)
Track lock min	ms/op	TLMN: Track lock latency minimum (ms/op)
Track lock avg	ms/op	TLAV: Track lock latency average (ms/op)
Track lock max	ms/op	TLMX: Track lock latency maximum (ms/op)
Track lock rate	ops/s	TLCN: Track lock operation rate (ops/s)
Track access min	ms/op	TAMN: Begin track access latency minimum (ms/op)
Track access avg	ms/op	TAAV: Begin track access latency average (ms/op)
Track access max	ms/op	TAMX: Begin track access latency maximum (ms/op)

Metric	Unit Axis	Description
Track access rate	ops/s	TACN: Track access operation rate (ops/s)

1.6.9.4 MDisk group cache partition performance metrics

Metrics of the performance-object type *MDisk group cache partition* derived from the object type *SVC node*.

Abbildung 71 Screenshot: options dialog: MDisk group cache partition performance options

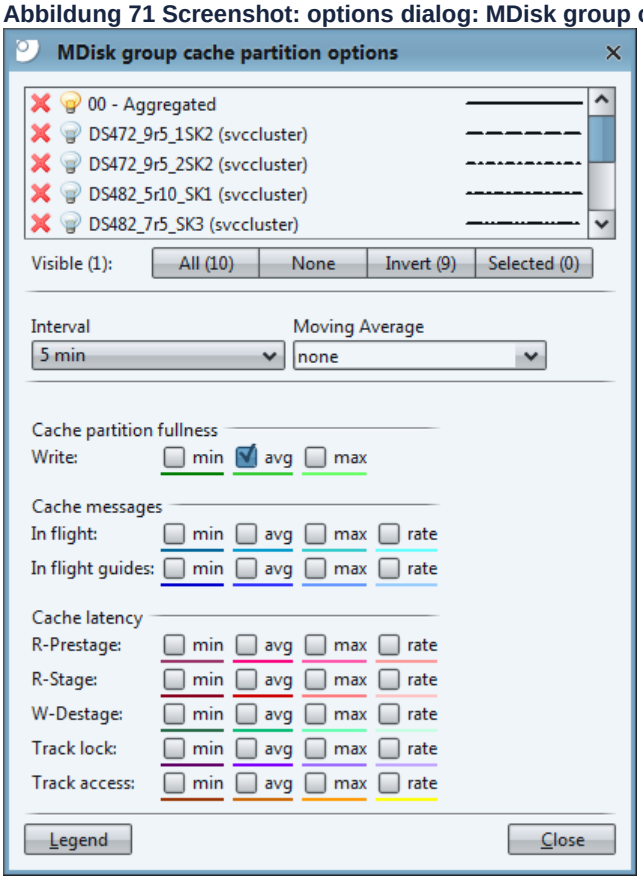

Tabelle 59 Overview: performance metrics MDisk group cache partition

Cache fullness				
Write:	✔ min	✔ avg	✔ max	
Cache messages				
In flight:	✔ min	✔ avg	✔ max	✔ rate

Cache fullness				
In flight guides:	✅ min	✅ avg	✅ max	✅ rate
Cache latency				
R-Prestage:	✅ min	✅ avg	✅ max	✅ rate
R-Stage:	✅ min	✅ avg	✅ max	✅ rate
W-Destage:	✅ min	✅ avg	✅ max	✅ rate
Track lock:	✅ min	✅ avg	✅ max	✅ rate
Track access:	✅ min	✅ avg	✅ max	✅ rate

Tabelle 60 Description: performance metrics MDisk group cache partition

Metric	Unit Axis	Description
Cache fullness		
W min	%	PFMN: Write cache fullness minimum (%)
W avg	%	PFAV: Write cache fullness average (%)
W max	%	PFMX: Write cache fullness maximum (%)
Messages		
In flight min	ops	MIFMN: Lowest messages in flight count in interval (ops)
In flight avg	ops	MIFAV: Average messages in flight count in interval (ops)
In flight max	ops	MIFNX: Highest messages in flight count in interval (ops)
In flight rate	ops/s	MIFCN: Message rate (count/s)
In flight guide min	ops	MIFGMN: Lowest messages in flight guide count in interval (ops)
In flight guide avg	ops	MIFGAV: Average messages in flight guide count in interval (ops)
In flight guide max	ops	MIFGNX: Highest messages in flight guide count in interval (ops)
In flight guide rate	ops/s	MIFGCN: Message guide rate (ops/s)

Metric	Unit Axis	Description
Cache latency		
W-Destage min	ms/op	DLMN: Destage latency minimum (ms/op)
W-Destage avg	ms/op	DLAV: Destage latency average (ms/op)
W-Destage max	ms/op	DLMX: Destage latency maximum (ms/op)
W_Destage rate	ops/s	DLCN: Destage rate (ops/s)
R-Prestage min	ms/op	PLMN: Prestage (read ahead) latency minimum (ms/op)
R-Prstage avg	ms/op	PLAV: Prestage (read ahead) latency average (ms/op)
R_Prestage max	ms/op	PLMX: Prestage (read ahead) latency maximum (ms/op)
R-Prestage rate	ops/s	PLCN: Prestage (read ahead) operation rate (ops/s)
R-Stage min	ms/op	SLMN: Stage (read miss) latency minimum (ms/op)
R-Stage avg	ms/op	SLAV: Stage (read miss)latency average (ms/op)
R-Stage max	ms/op	SLMX: Stage (read miss) latency maximum (ms/op)
R-Stage rate	ops/s	SLCN: Stage (read miss) rate (ops/s)
Track lock min	ms/op	TLMN: Track lock latency minimum (ms/op)
Track lock avg	ms/op	TLAV: Track lock latency average (ms/op)
Track lock max	ms/op	TLMX: Track lock latency maximum (ms/op)
Track lock rate	ops/s	TLCN: Track lock rate (ops/s)
Track access min	ms/op	TAMN: Begin track access latency minimum (ms/op)
Track access avg	ms/op	TAAV: Begin track access latency average (ms/op)
Track access max	ms/op	TAMX: Begin track access latency maximum (ms/op)
Track access rate	ops/s	TACN: Track access rate (ops/s)

1.6.9.5 Node port performance metrics

Metrics of the performance-object type *node port*, derived from the object type *SVC node*.

Abbildung 72 Screenshot: options dialog: node port perf options

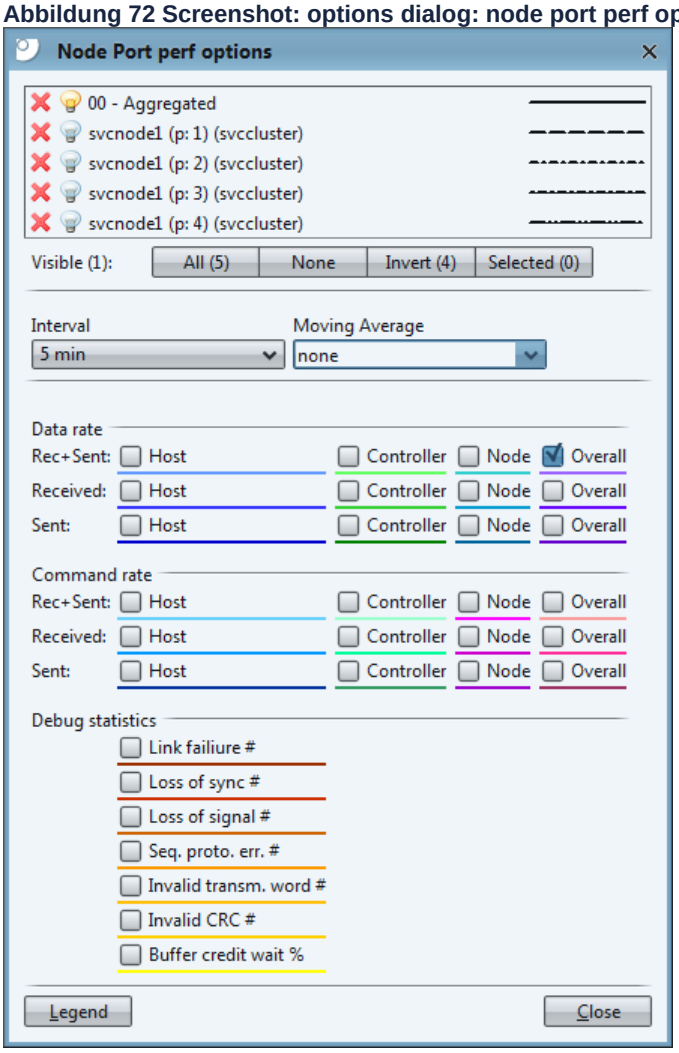

Tabelle 61 Overview: performance metrics node port

Debug statistics				
Data rate				
Rec + sent:	✅ Host	✅ Controller	✅ Node	✅ Overall
Received	✅ Host	✅ Controller	✅ Node	✅ Overall

Debug statistics				
Sent	✅ Host	✅ Controller	✅ Node	✅ Overall
Command rate				
Rec + sent:	✅ Host	✅ Controller	✅ Node	✅ Overall
Received	✅ Host	✅ Controller	✅ Node	✅ Overall
Sent	✅ Host	✅ Controller	✅ Node	✅ Overall
	✅ Link failiure #			
	✅ Loss of sync #			
	✅ Loss of signal #			
	✅ Seq. proto. err. #			
	✅ Invalid transmission word #			
	✅ Invalid CRC #			
	✅ Buffer credit wait %			

Tabelle 62 Description: performance metrics node port

Metric label	Unit Axis	Description
Data rate		
Rec + sent: host	MiB/s	HBR+HBT: host bytes received and transmitted(MiB/s)
Rec + sent: controller	MiB/s	CBR+CBT: controller bytes received and transmitted (MiB/s)
Rec + sent: node	MiB/s	LNBR+RNBR+LNBT+RNBT: local plus remote node bytes received and transmitted (MiB/s)
Rec + sent: overall	MiB/s	HBR+CBR+LNBR+RNBR + HBT+CBT+LNBT+RNBT: overall bytes received and transmitted(MiB/s)
Received: host	MiB/s	HBR: host bytes received (MiB/s)

Metric label	Unit Axis	Description
Received: controller	MiB/s	CBR: controller bytes received (MiB/s)
Received: node	MiB/s	LNBR+RNBR: local plus remote node bytes received (MiB/s)
Received: overall	MiB/s	HBR+CBR+LNBR+RNBR: overall bytes received (MiB/s)
Sent: host	MiB/s	HBR: host bytes sent (MiB/s)
Sent: controller	MiB/s	CBR: controller bytes sent (MiB/s)
Sent: node	MiB/s	LNBR+RNBR: local plus remote node bytes sent (MiB/s)
Sent: overall	MiB/s	HBR+CBR+LNBR+RNBR: overall bytes sent (MiB/s)
Command rate		
Rec + sent: host	IO/s	HER+HET: host events received and transmitted(IO/s)
Rec + sent:controller	IO/s	CER+CET: controller events received and transmitted (IO/s)
Rec + sent: node	IO/s	LNER+RNER+LNET+RNET: local plus remote node events received and transmitted (IO/s)
Rec + sent: overall	IO/s	HER+CER+LNER+RNER + HET+CET+LNET+RNET: overall events received and transmitted(IO/s)
Received: host	IO/s	HER: host events received (IO/s)
Received: controller	IO/s	CER: controller events received (IO/s)
Received: node	IO/s	LNER+RNER: local plus remote node events received (IO/s)
Received: overall	IO/s	HER+CER+LNER+RNER: overall events received (IO/s)
Sent: host	IO/s	HER: host events sent (IO/s)
Sent: controller	IO/s	CER: controller events sent (IO/s)
Sent: node	IO/s	LNER+RNER: local plus remote node events sent (IO/s)
Sent: overall	IO/s	HER+CER+LNER+RNER: overall events sent (IO/s)

Metric label	Unit Axis	Description
Debug statistics		
Link failiure #:	Count / interval	LF: link failiure count
Loss of sync #:	Count / interval	LSY: loss-of-synchronisation count
Loss of signal #:	Count / interval	LSI: loss-of-signal count
Seq. proto. err. #:	Count / interval	PSPE: primitive sequence protocol error count
Invalid transm. word #:	Count / interval	ITW: invalid transmission word count
Invalid CRC #:	Count / interval	ICRC: invalid CRC count
Buffer credit wait %:	%	BBCZ: % of interval time data to send was waiting for free switch buffer-to-buffer-credits (Critical when > 10%, B2B recommended min: 41 credits)

The metric **BBCZ** is relatively important for the buffer to buffer credit tuning of the SAN.
Recommendation: the metric should not be higher than 10% within an interval.
41 is an ideal B2B-number for the switch ports to which SVC ports are connected.

1.6.9.6 Node to node performance metrics

Metrics of the performance-object type *node to node*, are derived from the object type *SVC node*.

In the statistics group NODE_TO_NODE performance metrics between two nodes are gathered.

This statistic is crucial for assessing the connection to remote clusters via Metro/Global Mirror.

There are two ways to connect nodes:

1. Intra IOgroup traffic (cache coherence and intra-cluster Remote Mirror (rather rare case))
2. Intercluster traffic for Remote Mirror (Metro or Global Mirror)

Therefore, it is useful to group the metrics of connected nodes according to their IOgroup or their cluster partnership.

1. IOgroup partner node: node to node metrics of the IOgroup-partner node

2. Remote cluster nodes: SUM (node to node metrics of all nodes taken from one of four possible remote cluster)

Performance chart type "node to node"
NODE_TO_NODE statistics form a new performance chart type: "node to node"
For each combination of two intercommunicating nodes (node pairs) there is one object instance.

Example: with two two-node clusters, there are 4 * (4-1) = 12 possible node pairs in the object list of the panel option.

Abbildung 73 Screenshot: options dialog: node to node perf options

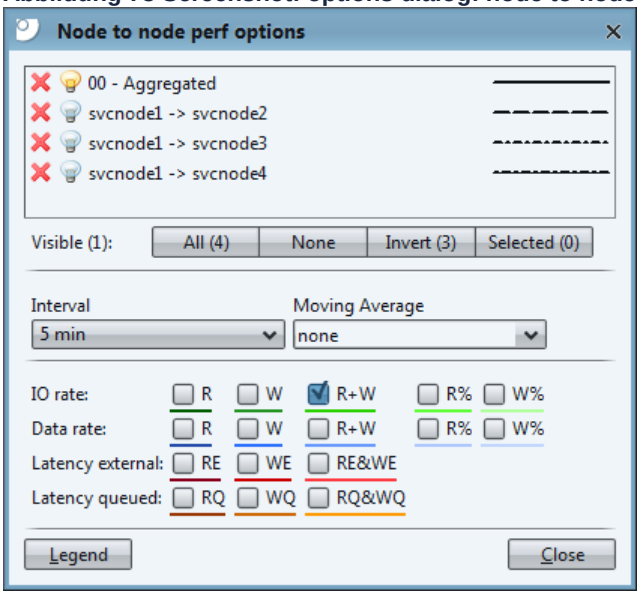

Tabelle 63 Overview: performance metrics node to node

Metrik					
IO rate:	✔ R	✔ W	✔ R+W	✔ R%	✔ W%
Data rate:	✔ R	✔ W	✔ R+W	✔ R%	✔ W%
Latency exten:	✔ RE	✔ WE	✔ RE&WE		
Latency queued:	✔ RQ	✔ WQ	✔ RQ&WQ		

Tabelle 64 Description: performance metrics node to node

Metric	Unit Axis	Description
IO rate R	ops/s	RO: Read operation rate (ops/s)
IO rate W	ops/s	WO: Write operation rate (ops/s)
IO rate R+W	ops/s	RO+WO: Total operation rate (ops/s)
IO rate R%	%	RO/(RO+WO): % Reads of total operations (%)
IO rate W%	%	WO/(RO+WO): % Writes of total operations (%)
Data rate R	MiB/s	RB: Read data rate (MiB/s)
Data rate W	MiB/s	WB: Write data rate (MiB/s)
Data rate R+W	MiB/s	RB+WB: Total data rate (MiB/s)
Data rate R%	%	RB/(RB+WB): % Read data of total data (%)
Data rate W%	%	WB/(RB+WB): % Write data of total data (%)
Latency RE	ms/op	RE/RO: Average response time excluding inbound queue time per read operation (ms/op)
Latency WE	ms/op	WE/WO: Average response time excluding outbound queue time per write operation (ms/op)
Latency RE&WE	ms/op	(RE+WE)/(RO+WO): Average response time excluding queue time per operation (ms/op)
Latency RQ	ms/op	RQ/RO: Average response time including inbound queue time per read operation (ms/op)
Latency WQ	ms/op	WQ/WO: Average response time including outbound queue time per write operation (ms/op)
Latency RQ&WQ	ms/op	(RQ+WQ)/(RO+WO): Average response time including queue time per operation (ms/op)

The metric **BBCZ** is important for the SAN´s buffer to buffer credit tuning.

Recommendation: the value should not exceed 10% within an interval.

The ideal B2B credit number on the switch ports connected to SVC ports is **41**.

1.7 Tables view

The tables view shows attributes of BVQ-relevant objects in tabular form.

Abbildung 74 Tables view: overview

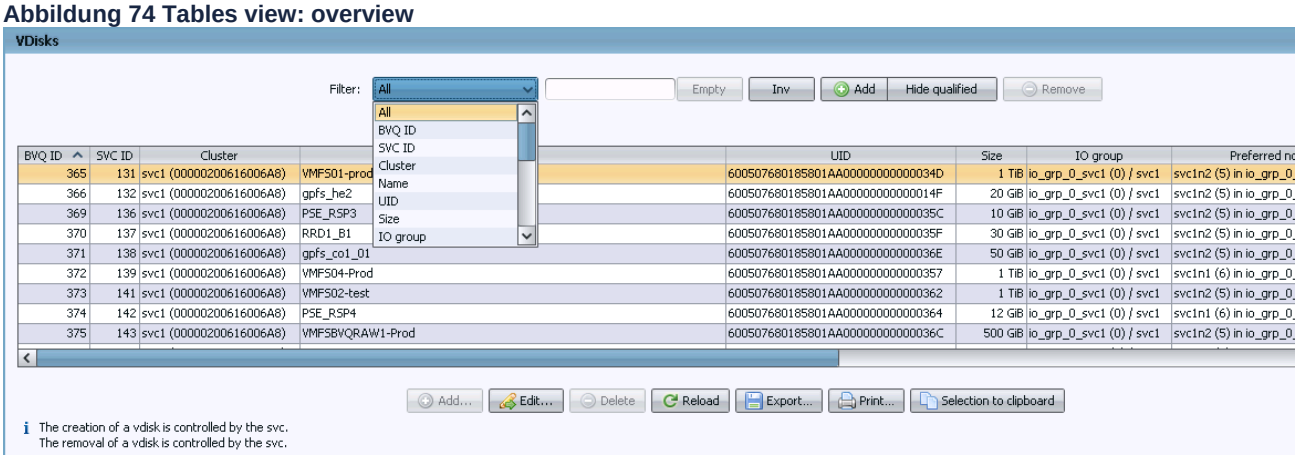

The tables view provides a tabular overview of all the objects of one type. Selected attributes are displayed in table columns.
The list can be sorted or filtered by attributes (columns). In addition, a targeted search is possible.

Table views do not exist for all objects but usually only for objects with object editors. For an overview, please see chapter object types.

Call of the tables view from the main menu:

Table views can be opened in the sub-menu "object-lists" in the main menu. Object lists are also available via the menu item "show in list" in context menus of BVQ objects.

Call from the treemap:

The list view can also be accessed from the topology view via "show in list" in the tree map context menu.
The selected objects are preselected in the list. They can be exported there.

Abbildung 75 Tables view: call via the treemap

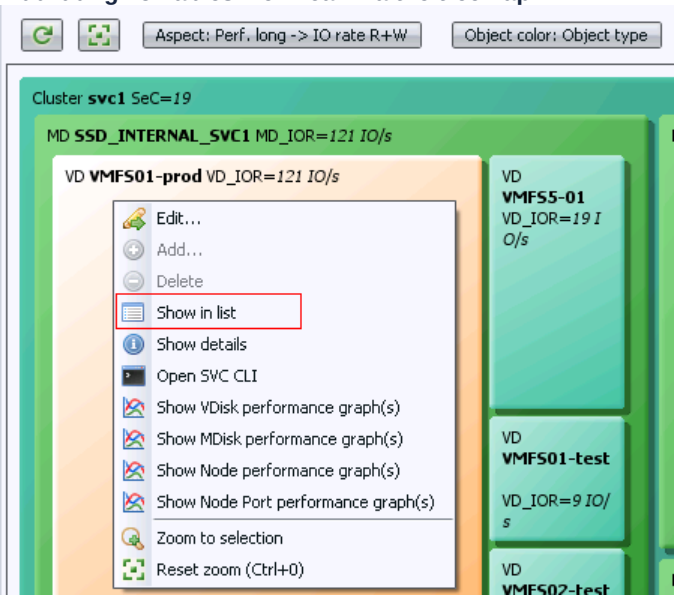

1.7.1 Filtering

Abbildung 76 Tables view: filter dialog

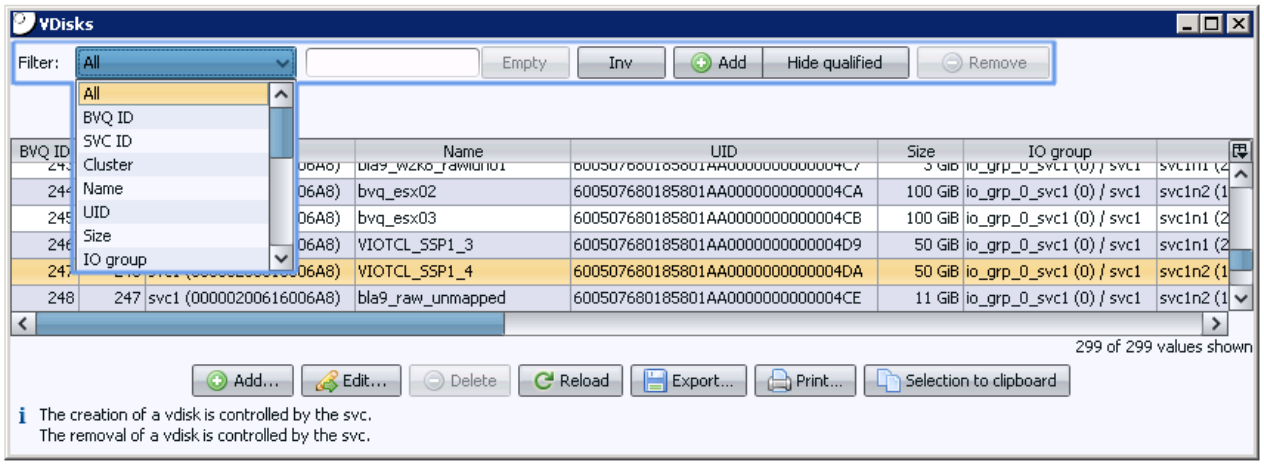

The filter in the table view allows to display or hide rows according to certain criteria. Therefore a string or number hast to be entered into the input field, which is compared to the selected object attribute.
The special attribute "all" is the default value, which is the representative of all columns.

- Usually, text-based attributes are compared in the mode "string-part", i.e. search results need to contain only one part of the string in the input field. The upper / lower case doesn´t matter.

In this mode, the search result is updated after every newy key stroke , making the desired rows quick and interactive to find.

- Number-based attributes are compared in the mode *exact*. Attributes must exactly match the entered numeric value.
  In this mode, the comparison is started after the entry´s confirmation by [return].

A table's view-filter has the following controls:

Tabelle 65 Tables view: filter dialog controls

Attributes	Limits the search to the selected attribute.
Text field	Input field for the comparison value in the form of a string (possibly with wildcard characters) or numeric values.
Empty	Allows to search for empty attributes. This function is not supported for the attribute selection "all".
Inv	Inverts the result set of the currently selected filter. So, only those objects for which the filter expression does not apply are displayed.
Add	Adds an additional filter, which is "AND"-linked to the already existing filters.
Hide qualified	Adds a filter which excludes objects in the state *BVQ-qualified*.
Remove	Removes the filter

The following wildcard characters can be used in the filter text:

Tabelle 66 Tables view: wildcards filter

[*]	Any characters
[?]	A single character

1.7.2 Sorting the table

Tabelle 67 Tables view: columns´ functions

Sorting	By default the data records are sorted according to dates and times. This sortation can be inverted by clicking on the column header.
Changing order	The column order can be adjusted by "drag and drop" on the column´s header.
Fading	Hide and show individual columns

in/out columns	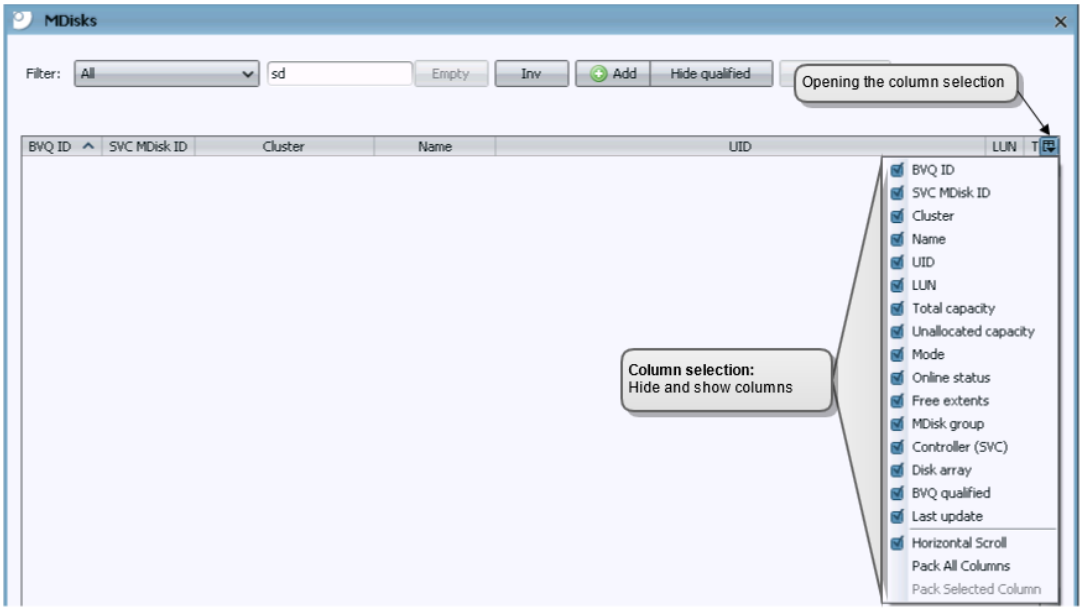
Adjusting column width	Allows change the column width either for keeping all column viewable or for making all column content viewable • **Horizontal scroll** Enables/disables the horizontal scrolls. When disabled, all columns are packed into the window by reducing the columns´ width. • **Pack all columns** The column width is optimized for all columns. • **Pack selected column** The column width is optimized for the selected column.

1.7.3 Search

The filter function described above can be augmented with the search window. The search window can be opened by using the key combination **[Ctrl + F]**.

Abbildung 77 Tables view: dialog

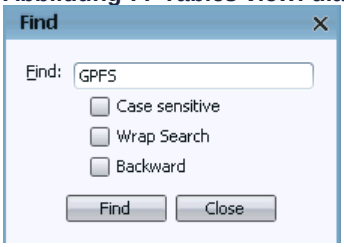

Controls:

Tabelle 68 Tables view: controls of the search dialog

Case sensitive	Activates the case-sensitivity for the search
Wrap	Search function starts again from the beginning after reaching the last entry
Backward	Search from bottom to top instead from top to bottom

A text search with wildcards is not supported here.

1.7.4 Other functions

By using the object's context menu (right click on a row) or the button row under the table, the following functions can be called:

Tabelle 69 Tables view: context menu´s functions

Add	Adds new objects in the with the object editor
Edit	Edits the attributes of the selected object in its object editor
Delete	Deletes the selected object
Show details	Calls the object detail view of the selected object
Reload	Reload the table from the GUI BVQ memory
Export	Exports the table into a CSV file (e.g. to import into MS Excel)
Print	Table printing
Selection to clipboard	Copies the selected items to the clipboard
Show ... performance graphs	Calls the performance view

Below the button line additional information is displayed if necessary. For example: if the object is already referenced by other objects, and therefore no deletion is possible.

1.8 Object detail view

The detail view represents detailed information about a single object in the form of a report.

This chapter is organized as follows:

- Calling the object detail view

- Overview object detail view
- Categories
- Information groups

1.8.1 Calling the object detail view

The object detail view can be accessed by **double-click** or **show details** from the context menu for a treemap object, the object list of the performance options dialogs as well as from the tables view.

Abbildung 78 Calling the object detail view

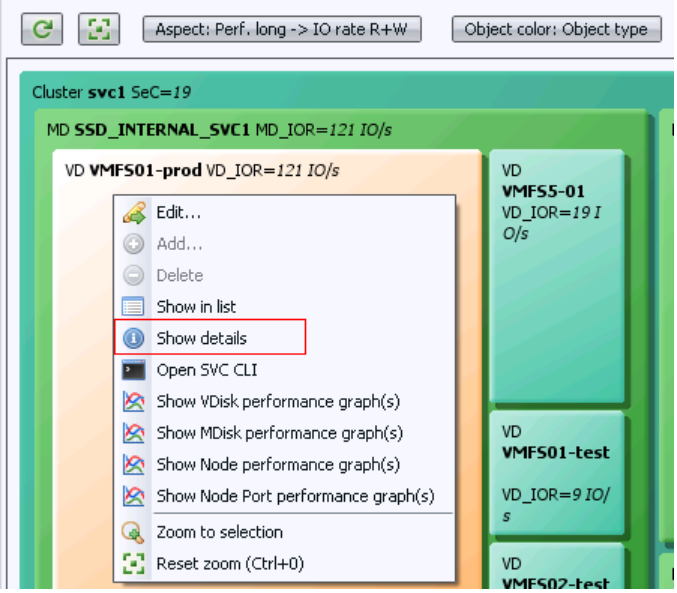

1.8.2 Overview object detail view

An object detail view is subdivided as follows:

- Categories (fold-out areas)
    - Subdivisions (tables with separate headings)
        - Information groups (common color markings)

1.8.3 Categories

Usually the object detail view is divided into the following main categories. In individually folding areas, corresponding information is displayed in tabular form.

- Attributes
- Accountables

- Topology
- Service level & compliance
- Performance

Abbildung 79 Object detail view: categories

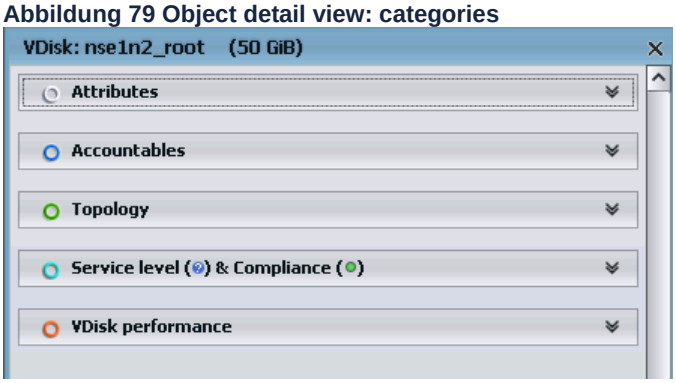

Depending on which topology - aspects is set, a matching category is already opened.

1.8.3.1 Attributes

This section contains object-specific information in three subdivisions:

Tabelle 70 Object detail view - categorie attributes: subdivisions

ID	Information needed to identify the respective component within the SVC and BVQ
SVC	SVC specific information / attributes
BVQ	BVQ specific additional information / attributes

If the object detail view is used, the section attributes is already opened in the aspect topology.

Abbildung 80 Object detail view: category attributes of a VDisk

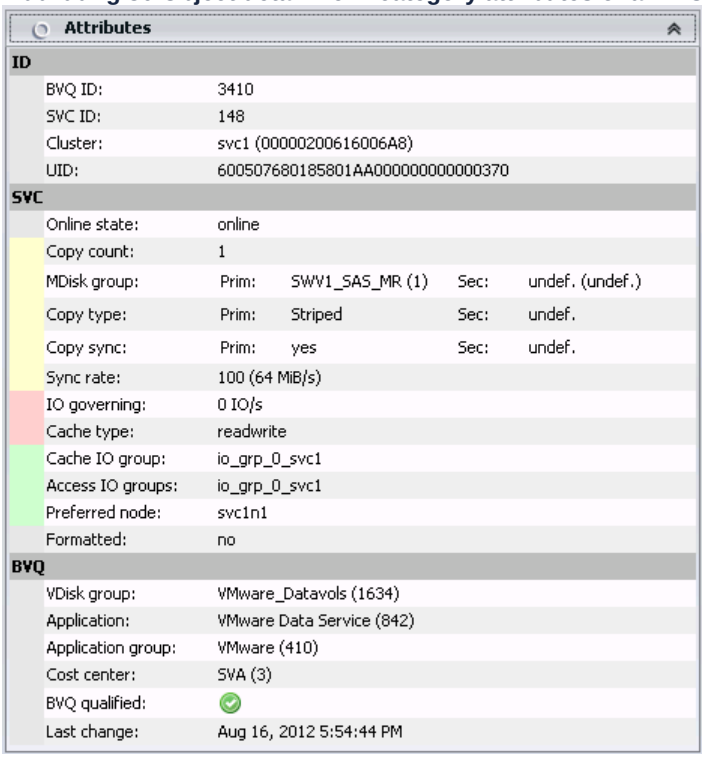

1.8.3.2 Accountables

This section summarizes billing factors in two subdivisions:

Tabelle 71 Object detail view - categoriy accountables: subdivisions

Capacities	Storage capacities – Key performance indicators for the management of storage capacities, such as assigned, free or virtual capacity of an object. –
Financials	Storage costs and prices – Financial indicators to support the charging and billing, such as storage costs and prices. –

It is also already shown in the aspekt capacity and the aspect cost & price, when the object detail view is opened.

Abbildung 81 Object detail view: category accountables of a MDisk group

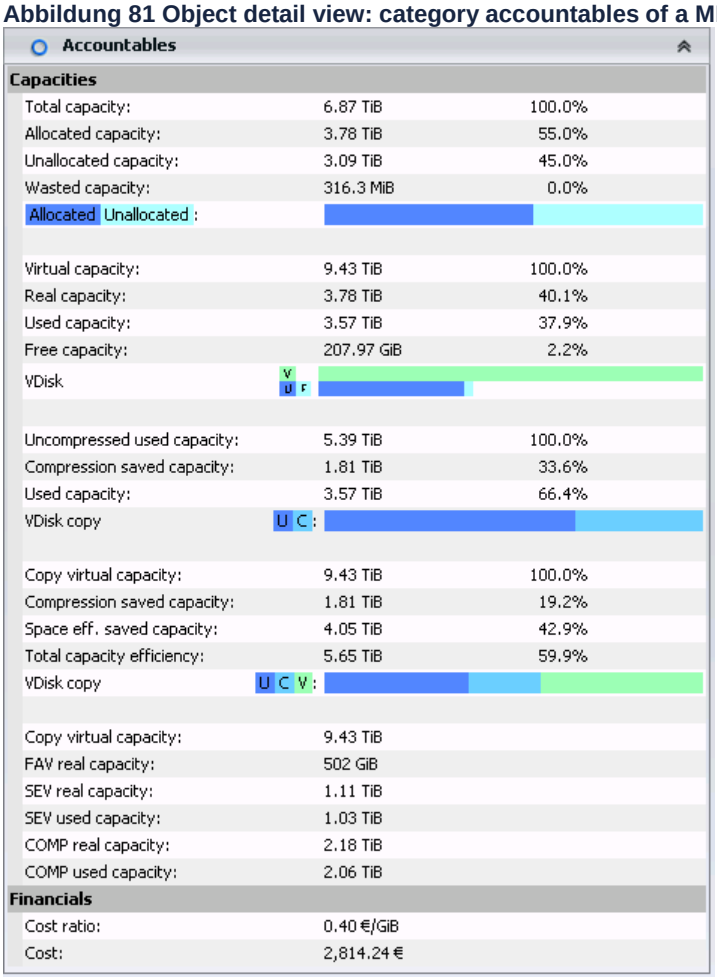
Accountables
Capacities
Total capacity: 6.87 TiB 100.0%
Allocated capacity: 3.78 TiB 55.0%
Unallocated capacity: 3.09 TiB 45.0%
Wasted capacity: 316.3 MiB 0.0%
Allocated Unallocated :
Virtual capacity: 9.43 TiB 100.0%
Real capacity: 3.78 TiB 40.1%
Used capacity: 3.57 TiB 37.9%
Free capacity: 207.97 GiB 2.2%
VDisk
V
u r
Uncompressed used capacity: 5.39 TiB 100.0%
Compression saved capacity: 1.81 TiB 33.6%
Used capacity: 3.57 TiB 66.4%
VDisk copy
U C :
Copy virtual capacity: 9.43 TiB 100.0%
Compression saved capacity: 1.81 TiB 19.2%
Space eff. saved capacity: 4.05 TiB 42.9%
Total capacity efficiency: 5.65 TiB 59.9%
VDisk copy
U C V :
Copy virtual capacity: 9.43 TiB
FAV real capacity: 502 GiB
SEV real capacity: 1.11 TiB
SEV used capacity: 1.03 TiB
COMP real capacity: 2.18 TiB
COMP used capacity: 2.06 TiB
Financials
Cost ratio: 0.40 €/GiB
Cost: 2,814.24 €

1.8.3.3 Topology

This are represents object counters – Number of objects linked with objects of another type..

It is already opened when the object detail view is activated in the aspect topology.

Abbildung 82 Object detail view: category topology of a MDisk group

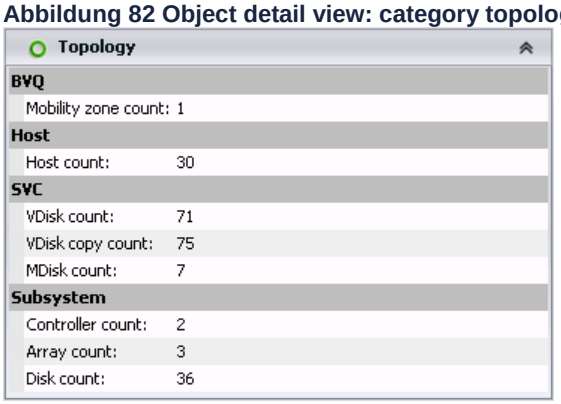

1.8.3.4 Service level & compliance

This section summarizes the results of the object qualification-rules´ evaluation in two subdivisions.

Tabelle 72 Object detail view - category service level & compliance: subdivisions

Service level	Contains descriptions for the assignment of the object into a higher grouping BVQ organisational unit and if necessary the currently valid BVQ SLA. **Placement rating:** Furthermore, the results of the evaluation of the configured BVQ DC to SC placement rating for the current BVQ data class / BVQ storage class combination. The following color schemes are designed similarly: color scheme data and storage class and color scheme DC2SC placement rating
Compliance	**Data class compliant:** contains the results from the above-named *service level*, but reduced to two states: 1. ✓ or ⚠ = ✓ and 2. ⊟ = ⊟ (color scheme data class compliance) **Mobility zone compliant:** contains the evaluation of the associated BVQ mobility zone placement rule (color scheme mobility zone compliance). **Vdisk group relation compliant:** contains the evaluation of the associated BVQ VDisk group relation placement rule (color scheme VDg relation compliance). **Overall compliant:** contains the overall aggregation-evaluation of the above-named compliance conditions (color scheme overall compliance)
Category heading	In the title the currently determined *placement rating* and *overall compliant* are already shown, so the category normally has not be opened.

Abbildung 83 Object detail view: category service level & compliance of a VDisk

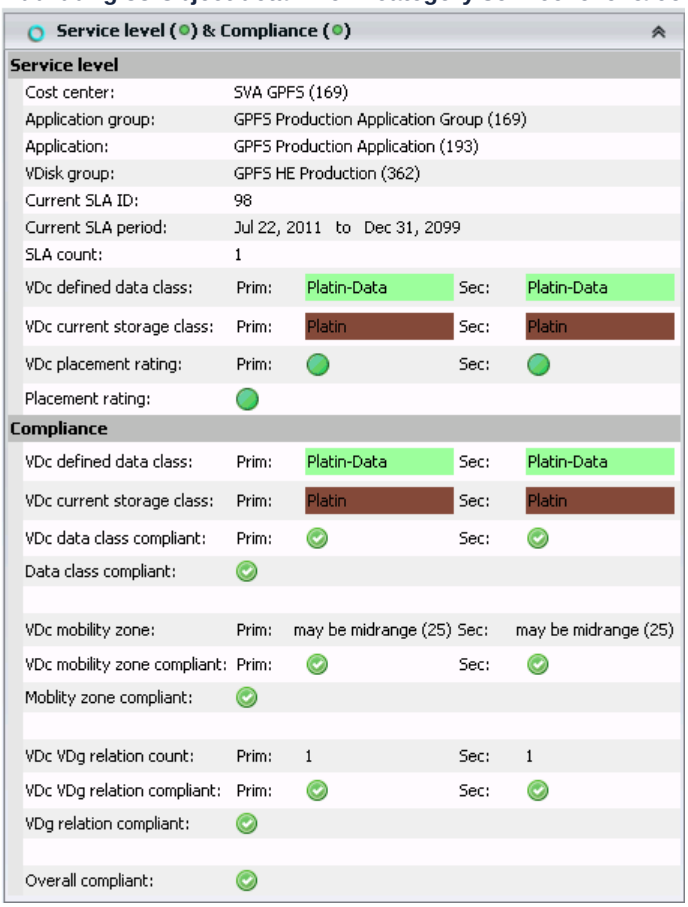
Service level (●) & Compliance (●)
Service level
Cost center: SVA GPFS (169)
Application group: GPFS Production Application Group (169)
Application: GPFS Production Application (193)
VDisk group: GPFS HE Production (362)
Current SLA ID: 98
Current SLA period: Jul 22, 2011 to Dec 31, 2099
SLA count: 1
VDc defined data class: Prim: Platin-Data Sec: Platin-Data
VDc current storage class: Prim: Platin Sec: Platin
VDc placement rating: Prim: Sec:
Placement rating:
Compliance
VDc defined data class: Prim: Platin-Data Sec: Platin-Data
VDc current storage class: Prim: Platin Sec: Platin
VDc data class compliant: Prim: Sec:
Data class compliant:
VDc mobility zone: Prim: may be midrange (25) Sec: may be midrange (25)
VDc mobility zone compliant: Prim: Sec:
Moblity zone compliant:
VDc VDg relation count: Prim: 1 Sec: 1
VDc VDg relation compliant: Prim: Sec:
VDg relation compliant:
Overall compliant:

1.8.3.5 Performance

This area is a tabular representation of the object´s performance indicators – Key figures for IO load, calculated mainly by averaging a series of measurements with different integration periods.

It is already displayed, when the object detail view is opened in the aspect performance.

Abbildung 84 Object detail view: category performance

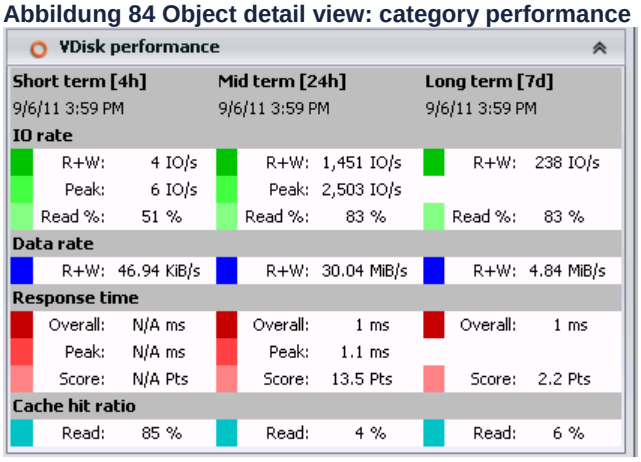

1.8.4 Information groups

If several information lines of a table are related to each other, it is indicated by the same color in the first narrow column.

Where possible, the following meanings of the colors were maintained:

Tabelle 73 Object detail view: colors of the information groups

Color	Used for
Light grey	No grouping
Light blue	Capacities or business objects
Light red	Performance
Light green	Related object links or topology objects
Light yellow	Other groupings
Light turquoise	Previously unused
Light orange	Previously unused

1.9 Key figures of the BVQ storage environment

Read and calculated key figures of the storage environment, managed by BVQ, about capacity, IO load and object counters.

Storage capacities	Key performance indicators for the management of storage capacities, such as assigned, free or virtual capacity of an object.
Storage costs and prices	Financial indicators to support the charging and billing, such as storage costs and prices.
Object counter	Number of objects linked with objects of another type.
Performance indicators	Key figures for IO load, calculated mainly by averaging a series of measurements with different integration periods.

1.9.1 Storage capacities

Key performance indicators for the management of storage capacities, such as assigned, free or virtual capacity of an object.

1.9.1.1 Visualization of the GUI´s storage capacities

Information about storage capacities are displayed in the GUI at four locations:

1. As a proportionality factor in the topology view's treemap in the aspect capacity
2. In the tooltip of the topology view´s treemap objects in the aspect group capacity
3. In the object detail view in the main chapter *accountables*
4. In the treemap color scheme group: capacity efficiency the degree of data reduction by compression and space efficiency is represented.

1.9.1.2 Metrics for storage capacities

Capacity related metrics and their aggregation logic are distinguished in two groups of object types:

Storage provider	Provide available capacity which is assigned to the storage consumers.
Storage consumer	Seize capacities which are provided by the storage providers.

D irect metrics of the storage providers

The following capacities metrics are available for **storage providers**:

TOTAL capacity	Total capacity *
ALLOCATABLE capacity	Assignable capacity *
ALLOCATED capacity	Assigned or occupied capacity
UNALLOCATED capacity	Unassigned capacity

* Difference: TOTAL – ALLOCATABLE

> *Under certain circumstances, the SVC/V7000 can not provide the total capacity of an object. E.g. if the MDISK-size is divided by its extent-size and doesn´t result in an integer.*

Direct metrics of the storage consumers

The following capacities metrics are available for **storage consumers**:

ALLOCATED capacity	Assignable capacity
VIRTUAL capacity	Virtual capacity *
REAL capacity	Space efficient VDisk copy real capacity **
USED capacity	Space efficient VDisk copy used capacity ***
UNCOMPRESSED_USED capacity	Space efficient VDisk copy uncompressd used capacity
FREE capacity	Space efficient VDisk copy free capacity ****
WASTED capacity	Differece between: ALLOCATED - REAL capacity – loss by only partially occupied extents

* **Virtual capacity:** Disk size which is presented to the host, and thus represents the maximum size of the storage consuming unit.

** **Real capacity:** Capacity which is initially allocated to the memory consuming unit at creation. If the VDisk parameter autoexpand is set to "yes", an extension is performed automatically when reaching USED = REAL capacity. This continues until REAL capacity = VIRTUAL capacity.

*** **Used capacity:** Part of the REAL capacity which was automatically assigned by host writes. The USED capacity always continues to grow in defined steps (grain size), until the REAL capacity is reached. The USED capacity can never be greater than the REAL capacity.

**** **Free capacity:** Capacity specification which is calculated from the difference between REAL capacity and USED capacity. The value indicates the free space within the REAL capacity, until its limit is reached and an extension of the REAL capacity is performed (if auto expand = yes).

Capacity rules per storage consumer type

Fully allocated VDisk (FAV): VIRTUAL capacity = REAL capacity = USED capacity / FREE capacity always 0

Space efficient VDisk (SEV): VIRTUAL capacity >= REAL capacity >= USED capacity / FREE capacity = REAL capacity - USED capacity

Fully allocated VDisk with VDisk mirror (FAV VDM): VIRTUAL capacity = 2 x REAL capacity = 2 x USED capacity / FREE capacity always 0

No inflexible rule. At the beginning the REAL and USED capacity can be less than the VIRTUAL capacity.
If REAL and USED capacity are growing because of write procedures, eventually the situation occurs, that the REAL and USED capacity become greater than the VIRTUAL capacity .
Finally, the state of a FAV VDisk copy can be reached.

Calculated complementary metrics

Consumer-metrics in **provider**-objects:
BVQ also calculates consumer-metrics for a *provider-object* by aggregating the values of linked consumer-objects.

Provider-metrics in **consumer**-objects:
In order ito get comparable values in the treemap visualization, **provider-metrics** are also assigned to *storage consumers* as follows:

TOTAL capacity	:=	REAL capacity
UNALLOCATED capacity	:=	0

1.9.1.3 Output of metrics in the object detail view´s information groups

In the object detail view, category "accountables", subdivision "capacities", the metrics are clearly divided and displayed in the following information groups:

Information group: storage pool

Capacity	Provider	Consumer	Min. granularity
TOTAL	✔	:= REAL	512 byte
ALLOCATABLE	✔	:= REAL	Extent
UNALLOCATED	ALLOCATABLE-∑	:= 0	Extent
ALLOCATED	∑	✔	512 byte
WASTED	∑	✔	512 byte
Bar chart: **allocated | unallocated**:	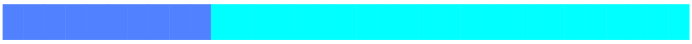		

Legend:

✔	Object´s native metric
∑	Calculated by aggregation of related objects

Information group: space efficiency (Thin Provisioning)

Capacity	Provider	Consumer	Min. granularity
VIRTUAL	∑	✔	512 Byte
REAL	∑	✔	512 Byte
USED	∑	✔	Grain size
FREE	∑	✔	Grain size
Bar chart: **V**irtual **U**sed **F**ree	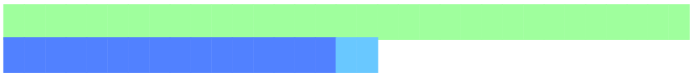		

Legend:

✔	Object´s native metric
∑	Calculated by aggregation of related objects

Information group: compression

Metric	Calculation	Description
Uncompressed used capacity:	= uncompressed_used_capacity	Theoretical size of the uncompressed used capacity
Compression saved capacity (Cs):	= uncompressed_used_capacity - used_capacity	Compression gain: saved capacity by compression
Used capacity (U):	= used_capacity	Compressed capacity
Bar chart: U Cs		Ratio of compressed (U) capacity to compression gain (Cs)

All percentages are based on the theoretical size of the uncompressed capacity (uncompressed used capacity = 100%).

Information group: capacity efficiency

Metric	Calculation		Description
Copy virtual capacity (V)	copy_virtual_capacity	= SUM(vdiskcopy.virtual_capacity)	Sum of the virtual capacity of all VDisk copies

Compression saved capacity:	compression_efficiency	= compression_saved_capacity	Compression gain: saved capacity by compression
Space eff. saved capacity:	space_efficiency	= copy_virtual_capacity - uncompressed_used_capacity	Reduction gain: saved capacity by Thin Provisioning
Total capacity efficiency	capacity_efficiency	= copy_virtual_capacity - allocated_capacity	Total saved capacity which is available for other usage in the pool
Bar chart: U Cs V	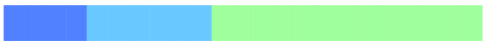		Ratio of compressed (U) capacity to compression gain (Cs) regarding the virtual capacity

All percentages are based on the sum of the virtual capacity of all VDisk copies (copy virtual capacity = 100%).

Note the difference between this simply calculated virtual capacity and the total virtual capacity:

- The simply calculated virtual capacity is the size of the volume which is visible from a server´s point of view.
- The total virtual capacity is calculated from the sum of all virtual copies' capacities of a volume and represents the comparison value for the identification of saved capacity

Simple example: A volume with a virtual capacity of 100GiB has 2 copies.

- Without Thin Provisioning or compression 200GiB of memory would have to be occupied.
- If the compression is activated for both copies and both achieve a compression rate of 50%, so 2 * 50 = 100GiB can be saved compared to the over the full assignment (200GiB).

Information group: type specific

For some objects´ subtypes, it makes sense to emit additional calculated metrics.

Especially for pool-creating storage providers like: SVC MDisk group, BVQ storage class, SVC cluster, BVQ mobility zone:

FAV REAL capacity	Sum of the fully allocated VDisk copy REAL capacity
SEV REAL capacity	Sum of the space efficient VDisk copy REAL capacity
SEV USED capacity	Sum of the space efficient VDisk copy USED capacity
COMP REAL capacity	Sum of the compressed VDisk copy REAL capacity

COMP USED capacity	Sum of the compressed VDisk copy USED capacity

For virtualization object types: SVC IO group und SVC node:

FOREIGN ALLOCATED capacity	Sum of the ALLOCATED capacity which is managed by the other SVC nodes or SVC IO groups of the same cluster

1.9.1.4 Aggregation of storage capacities

For object types without direct measurements, capacities are calculated by aggregating. (Please see: aggregation of storage capacities).

1.9.1.5 Aggregation of storage capacities

For object types without direct measurements, capacities are calculated by aggregating.

A direct calculation of storage capacities is only possible for the following object type metric combinations:

Object type	TOTAL	ALLOCATABLE	ALLOCATED	UNALLOCATED	VIRTUAL	REAL	USED	FREE
SVC extent	✓	0	✓	0	N/A	N/A	N/A	N/A
SVC MDisk	✓	✓	✓	✓	∑ VDisk	∑ VDisk copy	∑ VDisk copy	∑ VDisk copy
SVC MDisk group	✓	✓	✓	✓	∑ VDisk	∑ VDisk copy	∑ VDisk copy	∑ VDisk copy
SVC VDisk copy	✓	N/A	✓	N/A	= VDisk	✓	✓	✓
SVC VDisk	✓	N/A	∑ Extents	N/A	✓	∑ VDisk copy	∑ VDisk copy	∑ VDisk copy

For other object types capacity metrics are obtained by aggregation indirectly. Thus, for example, the VIRTUAL VDisk capacity for a VDisk group is calculated by summing up the VDisk group, which contains the combined VDisk VIRTUAL capacity. The grouping objects are used to select the

base objects. The group membership, which is displayed at this point of time, counts.

Aggregation methods

Aggregation method	Description
∑	The aggregated value is defined as the sum of (defined) basic values

Aggregated object types

Aggregated object type	TOTAL	ALLOCATABLE	ALLOCATED	UNALLOCATED	VIRTUAL	REAL	USED	FREE
BVQ controller	∑ MDisk	∑ MDisk	∑ MDisk	∑ MDisk	∑ VDisk	∑ MDisk	∑ MDisk	∑ MDisk
SVC controller	∑ MDisk	∑ MDisk	∑ MDisk	∑ MDisk	∑ VDisk	∑ MDisk	∑ MDisk	∑ MDisk
BVQ disk type	∑ MDisk	∑ MDisk	∑ MDisk	∑ MDisk	∑ VDisk	∑ MDisk	∑ MDisk	∑ MDisk
BVQ disk array	∑ MDisk	∑ MDisk	∑ MDisk	∑ MDisk	∑ VDisk	∑ MDisk	∑ MDisk	∑ MDisk
SVC controller port	∑ MDisk	∑ MDisk	∑ MDisk	∑ MDisk	∑ VDisk	∑ MDisk	∑ MDisk	∑ MDisk
SVC node	= Cluster	N/A	∑ VDisk	= Cluster	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
SVC IO group	= Cluster	N/A	∑ VDisk	= Cluster	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
SVC cluster	∑ MDisk	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ VDisk	∑ VDisk copy	∑ VDisk copy	∑ VDisk copy
BVQ VDisk group	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
BVQ applicatio	∑	∑ VDisk	∑ VDisk	∑ VDisk	∑	∑	∑	∑

Aggregated object type	TOTAL	ALLOCATABLE	ALLOCATED	UNALLOCATED	VIRTUAL	REAL	USED	FREE
n	VDisk				VDisk	VDisk	VDisk	VDisk
BVQ application group	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
BVQ cost center	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
SVC host	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
SVC VDisk to host map	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
BVQ data class	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
BVQ storage class	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ MDisk group
BVQ SLA	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
BVQ mobility zone	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ MDisk group	∑ VDisk	∑ MDisk group	∑ MDisk group	∑ MDisk group
BVQ VDisk group relation	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk	∑ VDisk
BVQ room	∑ Controller	∑ Controller	∑ Controller	∑ Controller	∑ VDisk	∑ Controller	∑ Controller	∑ Controller
BVQ site	∑ Controller	∑ Controller	∑ Controller	∑ Controller	∑ VDisk	∑ Controller	∑ Controller	∑ Controller

1.9.2 Storage costs and prices

Financial indicators to support the charging and billing, such as storage costs and prices.

1.9.2.1 Visualization in the GUI

Costs and prices are displayed in the GUI in three places:

1. As proportionality factor in the treemap of the topology view in the aspect cost & price
2. In the tooltip in treemap-objects of the topology view in the aspect cost & price
3. In the object detail view in the main chapter *accountables*

1.9.2.2 BVQ financial metrics

The allocation of costs and prices is equal for all items in BVQ:

- The **cost ratio** of the SVC MDisk group is multiplied by the *ALLOCATED capacity* of an object.
- The **price ratio** from the currently valid BVQ SLA of a VDisk group is multiplied by an object´s *VIRTUAL capacity*.

Tabular overview for the assignment of BVQ finances in memory provisioners- and memory consumers-object types:

Financial metrics	Factor	Provisioner	Consumer
COST	SVC MDisk group.Cost ratio	*TOTALcapacity*	*ALLOCATED capacity*
PRICE	BVQ SLA.Price ratio	N/A	*VIRTUAL capacity*

1.9.3 Object counter

Number of objects linked with objects of another type.

Object counter answer typical questions, such as:

- How many VDisks are in the cluster?
- How many spindles (HDDs) has my storage pool?
- How many hosts does the SVC cluster know?
- ...

1.9.3.1 Display in the GUI

Object counters are presented in the GUI in two places:

1. In the tooltip of topology view's treemap objects as subelement count (SeC)
2. In the object detail view in the main chapter *counter*

1.9.4 Performance indicators

Key figures for IO load, calculated mainly by averaging a series of measurements with different integration periods.

Performance Indicators allow a load comparison of selected objects in the topology view.

1.9.4.1 Visualization of performance indicators in the GUI

Performance indicators are displayed in the GUI at three places:

1. As proportionality factor in the treemap of the topology view in the aspect performance
2. In the tooltip of the treemap objects of the topology view in the aspect performance
3. In the object detail view in the main chapter *performance*

1.9.4.2 Integration periods

Performance indicators are calculated automatically for 4 different periods in the immediate past and recorded into the database. The calculation is performed regularly by a background pro cess.

The performance indicators are available in 4 different variants:

Indicator type	Description	Period	Calculation cycle	Description
Real time	Actual	5min	5min	Recent values of the actual past for direct response to the current situation at hand (e.g. performance bottleneck, load-top, ...)
Short term	Short-term	4h	1h	Prompt values of the recent past for an effective operational response to a longer enduring situation
Mid term	Medium-term	24h	6h	Medium-term values to match typical load cycles (e.g. last 24h)
Long term	Long-term	7d	24h	Long-term, "calmed" values to match typical load cycles (e.g. over the last 7 days)

The GUI always loads and displays the last calculated (current) indicators.

1.9.4.3 Performance indicator metrics

The following indicators are calculated based on the measured values of the relevant SVC objects (VDisk, MDISK, Node). Their definitions are identical to the respective metrics of the performance view.

For object types without direct measurement values, performance indicators are calculated by aggregation. (See: aggregation of performance indicators).

Indicator name	Shortcut	Calculation	VDisk	MDisk	Node	Note

Indicator name	Shortcut	Calculation	VDisk	MDisk	Node	Note
IO rate	IOR	IOs / time	✓	✓	–	
Peak IO rate	IORP	90-percentil of the 5min-IOR-samples	✓	✓	–	
Data rate	DR	Bytes / time	✓	✓	✓	
Peak data rate	DRP	90-percentil of the 5min-DR-samples	–	–	✓	
Response time / IO	RT	(cumulated resp. time / IO-quantity)	✓	✓	–	undefined for IOR < 10 IO/s
Response time score	RTS	(RT/10)^2 * IOR	✓	✓	–	undefined for IOR < 10 IO/s; *
Peak response time	RTP	90-percentil of the 5min-RT-samples	✓	✓	–	undefined for IOR < 10 IO/s
Read ratio	R%	Read IOs / IO quantity	✓	✓	–	
Read cache hit ratio	RCH%	Cache read hits / cache reads	✓	–	–	
CPU load	CPU	CPU-ticks / time	–	–	✓	1000 ticks per second, in %

* Response time score

> *RTS is a dimensionless rating number, which particularly highlights the originator of bottlenecks in the treemap.*

1.9.4.4 Aggregation of performance indicators

For object types without direct measurement values, performance indicators are calculated by aggregation.

A direct calculation of performance indicators (PI) only occurs for object types, which have measurement readings.
These PIs are calculated based for the basic object types: VDisk, MDisk, node.

For other object types PIs are obtained by aggregating the base-PIs (aggregated PIs). Thus, for example the PI "IO rate" for a VDisk group is calculated by summing the condensed VDisk PIs, which are condensed in the VDisk group.

- The grouping objects are used to select the base objects. The calculation uses the most recent state of the group-membership when aggregating.

- Undefined underlying data aggregation points are ignored. (note: for mean values correspondingly reduced number of elements)
- If more than 20% of the underlying data aggregation points are undefined, the aggregated PI is undefined.

Aggregation methods

Aggregation methods	Description
∑ (Sum)	The aggregated PI is defined as the sum of (defined) base-PIs
Ø (Mean)	The aggregated PI is defined as the arithmetic mean of the base-PIs (sum of the defined basic PIs / number of defined basic-PIs)

Performance indicator (Basic-PI)	Aggregation method
IOR	∑
DR	∑
RT/IO	Ø
RTS	Ø
R%	Ø
RCH%	Ø
CPU	Ø
xxxP	N/A *

* Peak-PIs are not defined for the aggregation.

Aggregated object types

Aggregated object type	Basic objekt type	Performance indicators
BVQ controller	MDisk	all
SVC controller	MDisk	all
BVQ disk type	MDisk	all
BVQ disk array	MDisk	all

Aggregated object type	Basic objekt type	Performance indicators
SVC controller port	MDisk	all
SVC IO group	Node	all
SVC cluster	Node	all
SVC MDisk group	MDisk	all
BVQ VDisk group	VDisk	all
BVQ application	VDisk	all
BVQ application group	VDisk	IOR, DR, RT, RT-Score, R%
BVQ cost center	VDisk	IOR, DR
SVC Host	VDisk	all
BVQ data class	VDisk	all
BVQ storage class	MDisk	all
BVQ SLA	VDisk	all
BVQ mobility zone	MDisk	all
BVQ VDisk group relation	VDisk	all

1.10 Object types

Overview of all object types available in BVQ

Abbildung 85 BVQ-known object types and their relations

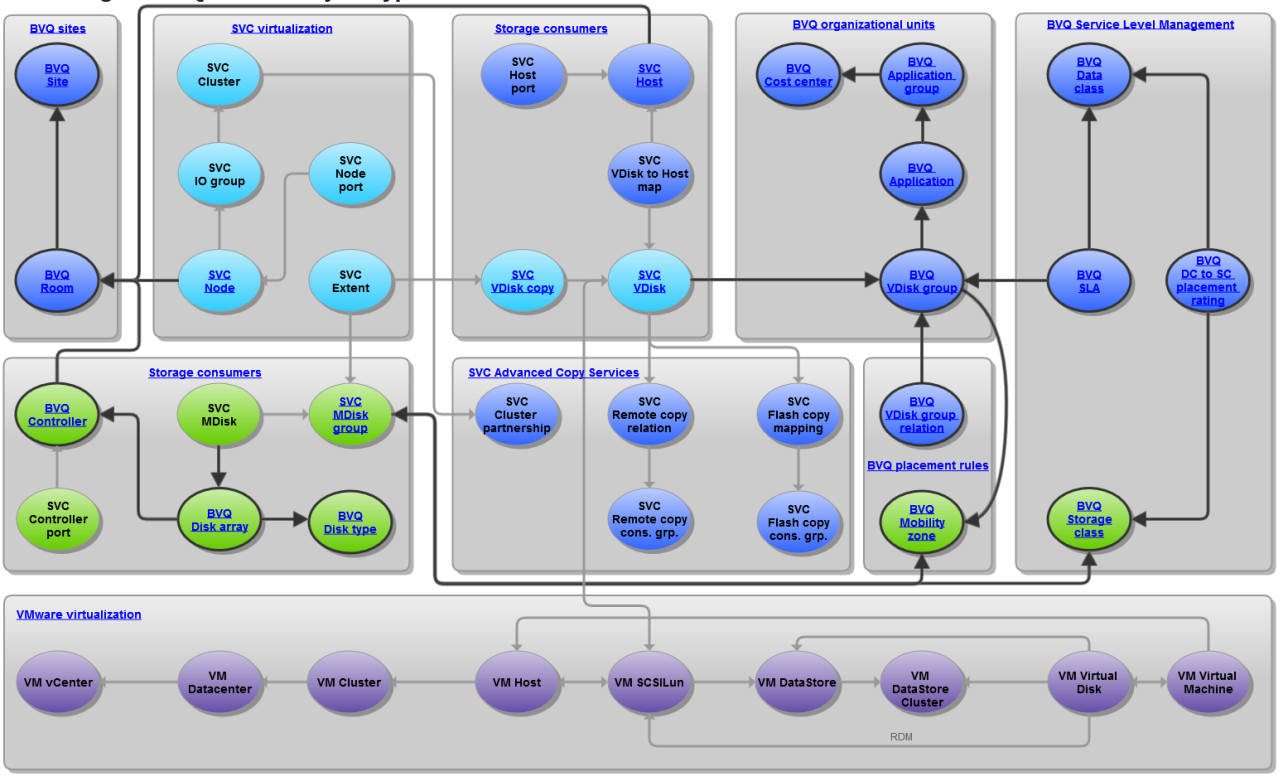

Caption:

Thin arrows	SVC internal object relations
Thick arrows	BVQ editable object relations Relation is configured in the source object
Thick edged objects	BVQ objects
Thin edged objects	SVC objects
Object colors	Objects´ standard colors in the BVQ GUI
Green objects	Objects of the IT-storage environment
Turquoise objects	Transform objects
Blue objects	Business-relevant objects

1.10.1 Overview object types

There are 2 different object types:

SVC objects	Representation of objects in the SVC BVQ database. Their attributes and capabilities are described in the IBM product documentation
BVQ objects	Management objects which are defined by BVQ and are associated with SVC objects

Main categories:

- Object-editors

Object editors
Not all objects are configurable in the BVQ GUI. These (SVC) objects don´t have any by BVQ modifiable attributes. So no object editors and no table views are existent.
Selected attributes of the other objects can be customized by the object editors of the BVQ GUI, or displayed in the tables view.

1.10.2 BVQ sites

Describes the locations of hardware like nodes, servers and storage subsystems.

BVQ room	Datacenter room within a site, which contains servers, storage subsystems and nodes.
BVQ site	Datacenter site containing one or more datacenter-operating rooms.

1.10.3 Storage providers

Objects that provide storage capacity (provider side).**Tabelle 74 Overview - storage providers**

BVQ disk type	Disk type found in disk arrays of storage subsystems.
BVQ disk array	Array (RAID) of one disk type in a storage subsystem.
BVQ controller	Represents a storage subsystem - cluster crossing group of SVC controllers.
SVC controller	Represents sub-units (port cards / controller-nodes with different WWNN) of a storage subsystem. *
SVC MDisk	Corresponds to a logical disk of a storage subsystem (LUNs), which is assigned to a SVC/V7000*
SVC MDisk group	Storage pool composed of SVC MDisk (logical disks)*

* For more details on the object type: See "IBM SVC software installation and configuration guide".

1.10.4 SVC virtualization

Units of SVC/V7000 virtualization layer.

SVC extent	Storage area which connects the SVC MDisks with SVC VDisk copies within a SVC MDisk group *
SVC node	SVC nodes or V7000 controller nodes.*
SVC IO group	Group of two SVC/V7000/V3700 nodes, which together provide the IO of a SVC VDisk´s subset *
SVC cluster	Groups all SVC objects belonging to one cluster. More than one SVC cluster can be managed by BVQ. *

* For more details on the object type: please see "IBM SVC software installation and configuration guide".

1.10.5 SVC advanced copy services

Advanced copy functions of the SVC, such as flash copy and remote copy.

SVC flash copy mapping	Flash copy mapping between source- and target SVC VDisk *
SVC flash copy consistency group	Consistency group of flash copy mappings *
SVC remote copy relationship	Remote copy relation between master- und auxiliary SVC VDisk *
SVC remote copy consistency group	Consistency group of remote copy relations *
SVC cluster partnership	Copy- partnership linking of two SVC clusters *

* For more details about the object type please see: "IBM SVC software installation and configuration guide

1.10.6 Storage consumers

Objects, space is allocated to (consumer side).

SVC VDisk copy	SVC VDisk copy corresponds to SVC object type VDisk copy. *

SVC VDisk	SVC/V7000 volume (virtual disk). *
SVC VDisk to host map	Allocation (mapping /LUN) of SVC VDisk to SVC host *
SVC host	Server which is connected to a SVC/7000 cluster and can be assigned to the SVC VDisks. *

* For more details about this object type: please see "IBM SVC software installation and configuration guide".

1.10.7 SAN connections and ports

Ports and connections of SAN devices, which are visible in SVC/V7000

SVC controller port	Port of a SVC controller *
SVC node port	Port of a SVC node *
SVC host port	Port of a SVC host *
SVC unknown port	Port of a SAN device with unknown type *
SVC port connect	Connection between a node port and another of the above SVC port types *

* For more details about this object type please see: "IBM SVC software installation and configuration guide"

Special color scheme:
SAN object types are always displayed in a special color scheme. Please see color scheme SVC controller port, SVC node port , SVC host port and color scheme port connect.

Source of these attributes:

Object type	Retrieves its attribute-values from the SVC CLI command *
SVC controller port	lscontroller (port details)
SVC node port	lsnode (port details)
SVC host port	lshost (port details)
SVC unknown port	lsfabric (filtered by type="unkonwn")
SVC port connect	lsfabric

* For more details about the SVC CLI commands please see: "IBM SVC command-line interface user's guide"

1.10.8 BVQ organizational units

Business units, business processes and departments.

BVQ VDisk group	Volume group - groups SVC VDisk of the same application and quality level.
BVQ application	Application (e.g. Oracle, Windows File Service, Oracle DB Service, SAP DB Service) grouping BVQ VDisk groups.
BVQ application group	Application group - groups applications e.g. for the representation of specialized procedures, business processes etc.
BVQ cost center	Cost center - combines application groups e.g. cost carriers, customers, and departments of a company.

1.10.9 BVQ placement rules

Placement guidelines for volumes

BVQ mobility zone	Group of storage pools to which a sub-set of assigned volumes can be placed.
BVQ VDisk group relation	Specifies a rule for the placement of two VDisk groups having the same usage context.

1.10.10 BVQ service level management

Storage and data classes, SLAs.

BVQ storage class	Storage class - group of storage pools with the same quality classification.
BVQ SLA	Service level agreement - associates a BVQ VDisk group with a BVQ data class and details on prices (€ / GiB) and run time.
BVQ data class	Data class - arranges volume-groups with equal quality requirements
BVQ DC to SC placement rating	Requirements for placement combinations of data- and storage class

1.10.11 BVQ - VMware virtualization

Storage-relevant object types of a VMware virtualization environment

VM vCenter server	VMware vCenter server
VM data center	VMware data center
VM cluster	VMware host cluster
VM host	VMware host
VM SCSI volume	VMware SCSI volume
VM datastore	VMware datastore
VM datastore cluster	VMware datastore cluster
VM virtual disk	VMware virtual disk
VM virtual machine	VMware virtual machine

For further details about these object types: please see VMware documentation.

1.10.12 BVQ organizational units

Business units, business processes and departments.

BVQ VDisk group	Volume group - groups SVC VDisk of the same application and quality level.
BVQ application	Application (e.g. Oracle, Windows File Service, Oracle DB Service, SAP DB Service) grouping BVQ VDisk groups.
BVQ application group	Application group - groups applications e.g. for the representation of specialized procedures, business processes etc.
BVQ cost center	Cost center - combines application groups e.g. cost carriers, customers, and departments of a company.

1.10.12.1 BVQ VDisk group

Volume group - groups SVC VDisk of the same application and quality level.

The "BVQ VDisk group" is a grouping object type, which combines one or more SVC VDisks.
Per BVQ VDisk group one or more period non-overlapping BVQ SLAs can be assigned with the connection to the BVQ data class (e.g. initially for a trial period, then for the production phase).

BVQ placement rules

- The freedom of movement of the associated SVC VDisks in SVC VDisk migration processes can be restricted by the assignment to a BVQ mobility zone.
- The placement of two VDisk groups to each other can be monitored by assigning this pair to a BVQ VDisk group relation.

It is useful for both methods, to divide the VDisk group according to their copy groups.
For example all primary-mirror VDisks into one (? _MP), the secondary-mirrors into the other VDisk group (? _MS).

Typical usage:

- Illustration of a group of VDisks with the same purpose (typically the same quality BVQ data class) and serving the same application BVQ application.
  e.g.: CD-images, DB tablespaces, DB redo log, filesystem volumes, archive volumes, system boot volumes.

Groups:

- SVC VDisk
- BVQ SLA
- BVQ VDisk group relation

Is grouped by: (configurable object relations)

- BVQ application
- BVQ mobility zone
- BVQ VDisk group relation

Configurable attributes

Name	Name of the VDisk group
Comment	Comment field

1.10.12.2 BVQ application

Application (e.g. Oracle, Windows File Service, Oracle DB Service, SAP DB Service) grouping BVQ VDisk groups.

The "BVQ application" is a grouping object type, which combines one or more BVQ VDisk groups.

Typcal usage:

- Illustration of applications referenced to a technical procedure, such as Oracle, File Service, Prod SAP, SAP Test

Groupes:

- BVQ VDisk group

Is grouped by: (configurable object relations)

- BVQ application group

Configurable attributes

Name	Name of the application
Comment	Comment field

1.10.12.3 BVQ application group

Application group - groups applications e.g. for the representation of specialized procedures, business processes etc.

The "BVQ application group" is a grouping object type, which combines one or more BVQ applications.

Typical usage:

- Representation of specialized procedures or business processes within a cost center or a cost carrier of a company
  e.g.: SAP HR, SAP Test, SAP development, SAP production.

Groups:

- BVQ application

Is grouped by: (configurable object relations)

- BVQ cost center

Configurable attributes:

Name	Name of the application group
Comment	Comment field

1.10.12.4 BVQ cost center

Cost center - combines application groups e.g. cost carriers, customers, and departments of a company.

The "BVQ cost center" is a grouping object, which combines one or more BVQ application groups.

Typical Usage:

- Representation of cost centers or cost units, customers, departments of a company, which subordinates one or more BVQ application groups.
  e.g.: human resources, end user name, SAP cost center 4711.

Groupes:

- BVQ application group

Is grouped by: (configurable object relations)

- No other type of objects

Configurable attributes

Name	Name of the cost center
Comment	Comment field

1.10.13 BVQ placement rules

Placement guidelines for volumes

BVQ mobility zone	Group of storage pools to which a sub-set of assigned volumes can be placed.
BVQ VDisk group relation	Specifies a rule for the placement of two VDisk groups having the same usage context.

1.10.13.1 BVQ mobility zone

Group of storage pools to which a sub-set of assigned volumes can be placed.

Application

SVC MDisk groups and BVQ VDisk groups are assigned to a mobility zone. Primarily to limit the placement´s decision for VDisks on selected MDisk groups, respectively to allow the examination if the current configuration complies with this placement rule.

SVC MDisk groups can be assigned to multiple mobility zones. This way they can be summarized according to different criteria.

BVQ VDisk groups are automatically subgrouped into VDisk copy type (P) primary or (S) secondary in VDisk copy groups. Each VDisk copy group may only be assigned to a single mobility zone.

A mobility zone serves ...

- ... the user to determine strict placement specifications.
- ... a possible algorithm for the selection of suitable MDisk groups for a VDisk migration.

- ... the compliance testing whether the current VDisk-placement corresponds to the standards. The result is displayed in the color scheme "mobility zone compliance" and in the object detail view (category "*service level and compliance*").

Application examples:

Mobility zones typically consist of:

- all MDisk groups of a controller
- all MDisk groups of a room or location
- all MDisk groups with the same disk- or RAID-type

The following image shows an example of a typical application with mirrored VDisks. In such a configuration of the mobility zone placement-rules, BVQ verifies their compliance and can warn the user in case of violations.

Abbildung 86 BVQ mobility zone: application example: two mirrored DB table VDisks shall be located in different sites

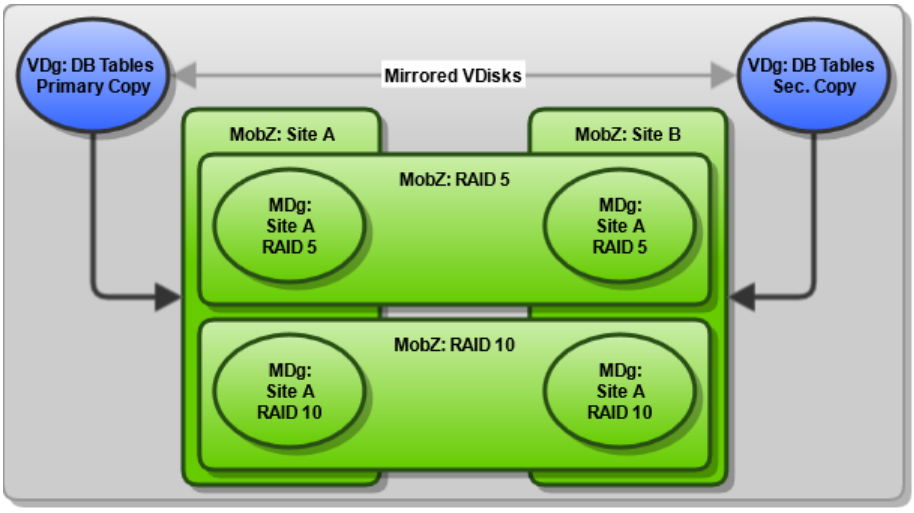
VDg: DB Tables Primary Copy
Mirrored VDisks
VDg: DB Tables Sec. Copy
MobZ: Site A
MobZ: Site B
MobZ: RAID 5
MDg: Site A RAID 5
MDg: Site A RAID 5
MobZ: RAID 10
MDg: Site A RAID 10
MDg: Site A RAID 10

Relations

Abbildung 87 BVQ mobility zone: configuration example

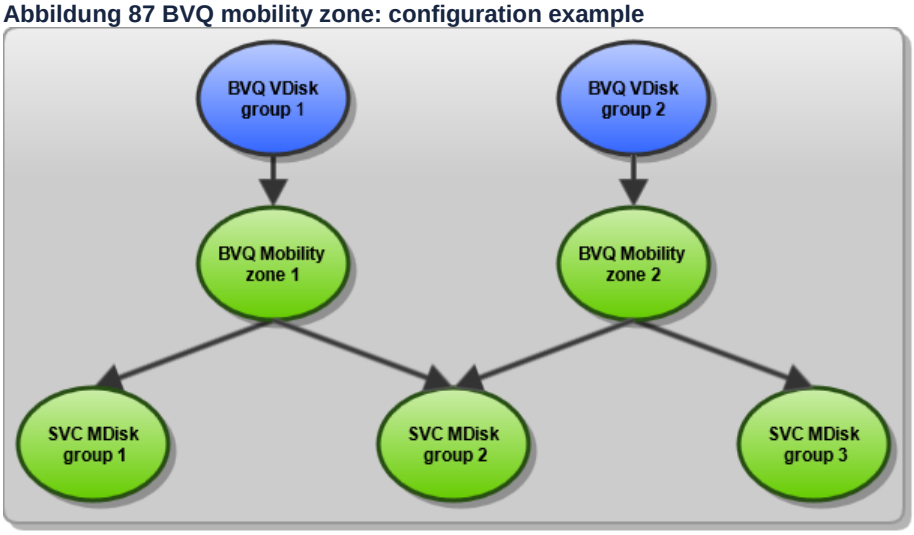
BVQ VDisk group 1
BVQ VDisk group 2
BVQ Mobility zone 1
BVQ Mobility zone 2
SVC MDisk group 1
SVC MDisk group 2
SVC MDisk group 3

Groups:

- BVQ VDisk group
- SVC MDisk group

Is grouped by: (configurable object relations)

- SVC MDisk group

Attributes

Name	Name of mobility zone
Comment	Comment field

Related object types

Compared to the BVQ VDisk group relation, which is also a placement rule, the BVQ mobility zone specifies the rule more conrete and establishes a direct relationship to properties (place, room, storage type, ..) of the potential target pool.

Their configuration requires less effort because the same BVQ mobility zone is potentially appropriate for many groups VDisk. Using the GUI´s mass edit feature is extremely helpful for configuration.

For rough allocation of VDisk groups to specific properties of a target pool, the BVQ mobility zone is the better choice. Once the placement has to be more detailed because of VDisk groups´ dependencies to each other, the BVQ VDisk group relation should be used additionally.

1.10.13.2 BVQ VDisk group relation

Specifies a rule for the placement of two VDisk groups having the same usage context.

Application example

Abbildung 88 BVQ VDisk group relation - application example: three VDisk groups in a complex context

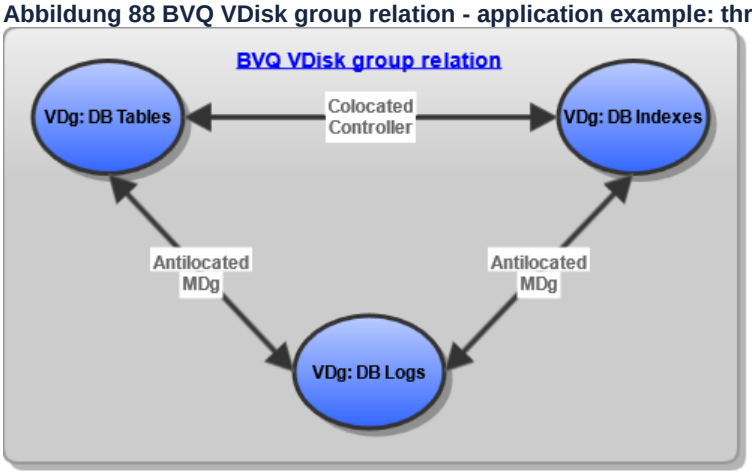

A VDisk group relation serves...

- ... the user to define a placeing rule of two BVQ VDisk groups, which are in an application context in the storage infrastructure.
  (For example: 2 copies of a VDisk mirror, 2 copies of a host LVM mirror, log, index and table data volumes of a database)
- ... a possible algorithm for the selection of suitable MDISK groups, storage systems, storage classes, ... in an automated VDisk migration
- ... the BVQ compliance check to verify if the current VDisk placement corresponds to the specifications which are displayed in the color scheme VDG relation compliance, in the category *service level & compliance* of the object detail view and in the treemap.

Relations

Abbildung 89 BVQ VDisk group relation: object type relations overview

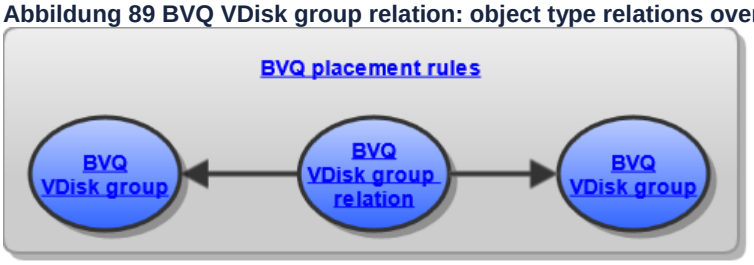

Groupes:

- A pair of two BVQ VDisk groups
  Several rules can be assigned to the same pair, if they don´t exclude each other.
  (For example: *colocated room* und *antilocated controller*, but not: *antilocated room* und *colocated controller*)

Is grouped by: (configurable object relations)

- BVQ VDisk group - Any number of VDisk group relations can be assigned to a VDisk group.

Attributes

Placement behavior	Selection of the condition: antilocated or colocated
Placement object type	Selection of the unit (object type), for which a condition is valid (e.g. anti-located site, colocated controller, ...)
VDdisk copy group (VDc)	When linking to a VDisk group it has to be specified whether it is a primary or a secondary group of the VDisk copies from the SVC / Storwize VDisk mirroring function :

GUI prefix	VDdisk copy group (VDc)
(P)	Subgroup of the **primary** VDisk copies
(S)	Subgroup of the **secondary** VDisk copies

Related object types

Compared to the BVQ mobility zone, which is also a placement rule, the BVQ VDisk group relation determines the rule more abstract and establishes no direct relation to properties (place, room, memory type, ..) of the potential target pools.

Their configuration requires more effort, since for each VDisk group exactly one fitting partner has to be found and there are potentially many more VDisk groups than BVQ mobility zones. A configuration with the mass edit feature of the GUI is not very helpful.

If only the raw allocation of VDisk groups to specific properties of target pools is considered, the BVQ mobility zone is the better choice. Once the placement has to be detailed further, due to dependencies of VDisk groups to each other, the BVQ VDisk group relation should be used as a supplement.

1.10.14 BVQ service level management

Storage and data classes, SLAs.

BVQ storage class	Storage class - group of storage pools with the same quality classification.
BVQ SLA	Service level agreement - associates a BVQ VDisk group with a BVQ data class and details on prices (€ / GiB) and run time.
BVQ data class	Data class - arranges volume-groups with equal quality requirements
BVQ DC to SC placement rating	Requirements for placement combinations of data- and storage class

1.10.14.1 BVQ storage class

Storage class - group of storage pools with the same quality classification.

Represents the actual classification of a group of storage pools and defines their operator-cost factor.
The attributes of the storage pool´s hardware (disk type, RAID Mode, quality and performance of the disk subsystem) typically represent the belonging to a storage class.
The storage class is closely related to the BVQ data class, which in contrast, represents the desired quality class of the managed virtual storage volumes (volumes group).

The "color" attribute is the basis for the staining of affected objects in the treemap in the color scheme data-, storage class.

Groups:

- SVC MDisk group
- BVQ DC to SC placement rating

Is grouped by: (configurable object relations)

- No other BVQ objects

Configurable attributes

Name	Name of the storage class

Color	Selection of the imaging color within the treemap
Default cost ratio	Default value for the storage costs in €/GiB allocated in assigning a SVC MDisk group.
Default design response time	Default value for the desired response time in ms/IO, allocated in assigning a SVC MDisk group.
Comment	Comment field

1.10.14.2 BVQ SLA

Service level agreement - associates a BVQ VDisk group with a BVQ data class and details on prices (€ / GiB) and run time.

A BVQ SLA belongs to exactly one BVQ VDisk group, but per BVQ VDisk group one or more BVQ SLA can be configured, which don´t overlap themselves in a time period. This also applies for the referenced objects, which are describing the service level. Changing and shortening are only possible by deleting and re-creating.

The BVQ SLA includes information about:

- Validity period (start and end date)
- With the end-customer agreed price factor for the virtual capacity of his SVC VDisk in €/GiB.

Groupes:

- No other objects

Is grouped by: (configurable object relations)

- BVQ data class
- BVQ VDisk group

Configurable attributes

Valid from	Date valid from ...
Valid to	Date valid to ...
Price ratio	End-customer price per GiB in € (versus costs of the SVC MDisk group)
Comment	Comment field

1.10.14.3 BVQ data class

Data class - arranges volume-groups with equal quality requirements

Represents the target classification of a volumes´ group and defines the retail price factor.
The description of a data class is typically found in service level performance-catalogs for end users.
It is closely related to the BVQ storage class, which is on the contrary the quality class of managed storage-hardware (group of storage pools).

The "color" attribute is the base for the staining of affected objects in the treemap with color scheme data-, storage class.

Groupes:

- Indirectly (via the currently valid BVQ SLA) BVQ VDisk group
- BVQ SLA
- BVQ DC to SC placement rating

Is grouped by: (configurable object relations)

- No other object type

Configurable attributes

Name	Name of the data class
Color	Selection of the representation color within the treemap
Default price ratio	Default value of the retail store price in €/GiB (versus default cost ratio in the BVQ storage class), which is allocatd to a BVQ SLA during the assignment.
Comment	Comment field

1.10.14.4 BVQ DC to SC placement rating

Requirements for placement combinations of data- and storage class

A "placement rating", contains a user configurable rating for each combination from data classes to storage classes .
This serves to propose a suitable placement of VDisks in MDISK groups (storage classes), and additionally to display an assessment of the actual situation the color scheme DC2SC placement rating.

Groupes:

- No other BVQ objects

Is grouped by: (Configurable object relations)

- BVQ storage class
- BVQ data class

Configurable attributes

<table>
<tr><td>Placement rating</td><td>Rating the placement for all combinations of data and storage classes.<table><tr><th>Symbol</th><th>Placement rating</th></tr><tr><td></td><td>Ideal placement</td></tr><tr><td></td><td>To risky or expensive placement</td></tr><tr><td></td><td>Not SLA compliant placement</td></tr></table></td></tr>
<tr><td>Data class</td><td>Arranges volume groups with the same quality requirements</td></tr>
<tr><td>Storage class</td><td>Group of storage pools with the same quality (classification)</td></tr>
</table>

1.10.15 BVQ sites

Describes the locations of hardware like nodes, servers and storage subsystems.

BVQ room	Datacenter room within a site, which contains servers, storage subsystems and nodes.
BVQ site	Datacenter site containing one or more datacenter-operating rooms.

1.10.15.1 BVQ room

Datacenter room within a site, which contains servers, storage subsystems and nodes.

Typical usage:

- Reports, for example available or allocated capacity per room
- Examination of DR compliance, e.g. for the question: are LVM mirrored volumes really spread over two rooms?
- Support for placement decisions, e.g. after VDiskGroup relation "antilocated room"

Groupes:

- SVC host
- SVC node
- BVQ controller

Is grouped by: (configurable object relationships)

- BVQ site

Configurable attributes

Name	Name of the room
Comment	Field for comments

1.10.15.2 BVQ site

Datacenter site containing one or more datacenter-operating rooms.

Typical usage:

- Reports, for example available or allocated capacity per site
- DR compliance testing , for example the question: Are LVM mirrored volumes actually distributed on two sites?
- Support for placement decisions, for example: according to the VDisk group relation "antilocated site"

In general, a site should be a completely self-contained functional unit which, in case of a disaster, can assume the most important IT services.

For this reason sites should be functionally independent of each other.

Groupes:

- BVQ room

Is grouped by: (configurable object relations)

- No other BVQ objects

Configurable attributes

Name	Site´s name
Comment	Comment field

1.10.16 BVQ - VMware virtualization

Storage-relevant object types of a VMware virtualization environment

VM vCenter server	VMware vCenter server
VM data center	VMware data center
VM cluster	VMware host cluster
VM host	VMware host

VM SCSI volume	VMware SCSI volume
VM datastore	VMware datastore
VM datastore cluster	VMware datastore cluster
VM virtual disk	VMware virtual disk
VM virtual machine	VMware virtual machine

For further details about these object types: please see VMware documentation.

1.10.17 SAN connections and ports

Ports and connections of SAN devices, which are visible in SVC/V7000

SVC controller port	Port of a SVC controller *
SVC node port	Port of a SVC node *
SVC host port	Port of a SVC host *
SVC unknown port	Port of a SAN device with unknown type *
SVC port connect	Connection between a node port and another of the above SVC port types *

* For more details about this object type please see: "IBM SVC software installation and configuration guide"

Special color scheme:
SAN object types are always displayed in a special color scheme. Please see color scheme SVC controller port, SVC node port , SVC host port and color scheme port connect.

Source of these attributes:

Object type	Retrieves its attribute-values from the SVC CLI command *
SVC controller port	lscontroller (port details)
SVC node port	lsnode (port details)
SVC host port	lshost (port details)
SVC unknown port	lsfabric (filtered by type="unkonwn")
SVC port connect	lsfabric

* For more details about the SVC CLI commands please see: "IBM SVC command-line interface user's guide"

1.10.18 Storage consumers

Objects, space is allocated to (consumer side).

SVC VDisk copy	SVC VDisk copy corresponds to SVC object type VDisk copy. *
SVC VDisk	SVC/V7000 volume (virtual disk). *
SVC VDisk to host map	Allocation (mapping /LUN) of SVC VDisk to SVC host *
SVC host	Server which is connected to a SVC/7000 cluster and can be assigned to the SVC VDisks. *

* For more details about this object type: please see "IBM SVC software installation and configuration guide".

1.10.18.1 SVC VDisk

SVC/V7000 volume (virtual disk).

Groups:

- SVC VDisk copy
- SVC VDisk to host map

Is grouped by: (configurable object attributes)

- BVQ VDisk group

Configurable attributes

No attributes

1.10.18.2 SVC Host

Server which is connected to a SVC/7000 cluster and can be assigned to the SVC VDisks.

Groupes:

- SVC VDisk to Host map
- SVC Host Port

Is grouped by: (configurable object relations)

- BVQ room

Configurable attributes

Comment	Comment field

1.10.18.3 SVC VDisk copy

SVC VDisk copy corresponds to SVC object type VDisk copy.

Groupes:

- Extents

Is grouped by: (configurable object relations)

- SVC VDisk

Configurable attributes

- No, object cannot be edited

The VDisk copy object is always a subobject of the VDisk.

If a VDisk mirror exits, then the two mirrors are grouped in VDisk.
If there is no VDisk mirror, then there is a 1:1 relation between VDisk and the associated VDisk copy object.

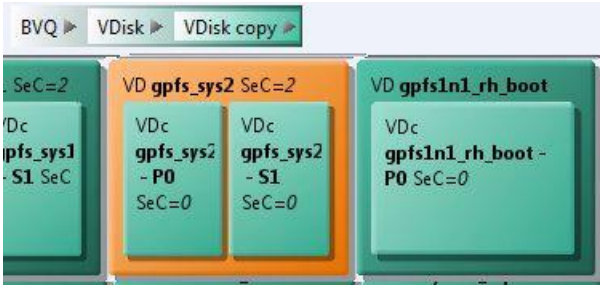

1.10.19 Storage providers

Objects that provide storage capacity (provider side).

Tabelle 75 Overview - storage providers

BVQ disk type	Disk type found in disk arrays of storage subsystems.
BVQ disk array	Array (RAID) of one disk type in a storage subsystem.
BVQ controller	Represents a storage subsystem - cluster crossing group of SVC controllers.
SVC controller	Represents sub-units (port cards / controller-nodes with different WWNN) of a storage subsystem. *

SVC MDisk	Corresponds to a logical disk of a storage subsystem (LUNs), which is assigned to a SVC/V7000*
SVC MDisk group	Storage pool composed of SVC MDisk (logical disks)*

* For more details on the object type: See "IBM SVC software installation and configuration guide".

1.10.19.1 BVQ disk type

Disk type found in disk arrays of storage subsystems.

Describes a type-class of disks which are integrated in the managed BVQ disk array.
Its technical attributes represent typical questions: how big?, how fast? and how reliable?

Groupes:

- BVQ disk array

Is grouped by: (configurable object relations)

- No other BVQ objects

Configurable attributes

Class	Classification of the disk´s type (enterprise / lowrange)
Size	Disk´s size (capacity) in GB (not in GiB)
Rotation speed	Rotation speed in U/min
Design IO rate	Upper limit of the IO-rate (IO/s) for operation under typical load profiles. Based on this value, the *estimated IO-density* is calculated in linked BVQ disk arrays.

1.10.19.2 BVQ disk array

Array (RAID) of one disk type in a storage subsystem.

Represents the attributes of a (RAID-) group of disk s within a storage system.
This is a delimited IO-activity unit.

Typical usage:

- RAID 5 array inside a storage system with 7 +1 plates, results in an IO effective disk count = 8
- Some systems combine internally several RAIDs. Here the total number of IO-effective disks has to be specified.

- Other systems (like IBM XIV) operate all drives in a single assemblage. These are entered as an array of all disks into BVQ (for example, in XIV R10).

Groupes:

- SVC MDisk

Is grouped by: (configurable object relations)

- BVQ controller
- BVQ disk type

Configurable attributes

Name	Name of the disk arrays
RAID level	Array´s RAID level (R1, R3, R4, R5, R6, R10, ...)
Disk type	Array´s BVQ Disk type
Disk count	Number of IO-effective plates (e.g., for a RAID5 with 7 +1 plates the value of 8 has to be entered).
Estimated IO density	Estimated or desired upper limit of IO-density (IOps / GiB) of the disk array. This is the base value from which IO-density of grouping storage provider-objects derived from. This is the threshold for the color scheme IO density heat.

More hints:

- RAID DP is registered as a RAID 6
- The IO-density of SSD drives can be rounded to 10 IOPS / GiB

Determination of IO-density for Easy Tier drives
A SSD has a assumed IO-density of 10IOPS/GiB, and a certain percentage of SSD storage is contained in our Easy Tier drive.
The IO-density of the disk part IODDisk is calculated with BVQ, and corrected by using the following formula:

$$IOD_{Hybrid} = (1-n)\ IOD_{Disk} + n^*\ IOD_{SSD}$$

Beispiel:

IO Dichte Disk:	IOD_{Disk}	= 0.5 IOPS/GiB
IO Dichte SSD:	IOD_{SSD}	= 10 IOPS/GiB
SSD Anteil	n	= 7% = 0.07

$$IOD_{Hybrid} = (1-0.07) * 0.5 + 0.07 * 10 = 1.17\ IOPS/GiB$$

1.10.19.3 BVQ controller

Represents a storage subsystem - cluster crossing group of SVC controllers.

A "BVQ controller" is a grouping object type, which combines one or more SVC controllers.

Typical usage:
Aggregation of one or more SVC controllers to a disk storage subsystem.
This is useful if:

- The SVC detects a storage subsystem not as one entity, but in several subunits (port cards / controller nodes with differing WWNN).
- A subsystem is connected to more than one SVC cluster.

Groupes:

- SVC controller
- BVQ disk array
- SVC MDisk

Is grouped by: (configurable object relations)

- BVQ room

Configurable attributes	
Name	Name of the storage subsystem
Class	Classification of the storage subsystem
Cache size	Size of the cache
MTM	Machine type and model - system description
SerNum	Serial number
Comment	Comment field

1.10.19.4 SVC controller

Represents sub-units (port cards / controller-nodes with different WWNN) of a storage subsystem.

A "SVC controller" is an image of the SVC object "controller", which represents sub-units (port cards / controller nodes with differing WWNN) of a storage subsystem.

Typical usage:

- Is combined in a BVQ controller to a disk storage subsystem.

Groupes:

- No other BVQ objects

Is grouped by: (configurable object relations)

- BVQ controller

Configurable attributes

- The SVC node can be associated to a BVQ controller.

1.10.19.5 SVC MDisk

Corresponds to a logical disk of a storage subsystem (LUNs), which is assigned to a SVC/V7000

Groups:

- No other BVQ objects

Is grouped by: (configurable object relations)

- BVQ disk array
- SVC controller
- SVC MDisk group

Configurable attributes

- No configurable attributes

1.10.19.6 SVC MDisk group

Storage pool composed of SVC MDisk (logical disks)

Groupes:

- SVC MDisk
- BVQ mobility zone
- SVC extent

Is grouped by: (configurable object relations)

- BVQ mobility zone
- BVQ storage class

Configurable attributes

Cost ratio	Storage costs per GiB in € (versus retail price in the BVQ SLA of the BVQ VDisk group)
Design response time	Maximum target response time of the storage pool based on its SVC MDisks. This is the threshold for the color scheme response time compliance.
Comment	Comment field

1.10.20 SVC advanced copy services

Advanced copy functions of the SVC, such as flash copy and remote copy.

SVC flash copy mapping	Flash copy mapping between source- and target SVC VDisk *
SVC flash copy consistency group	Consistency group of flash copy mappings *
SVC remote copy relationship	Remote copy relation between master- und auxiliary SVC VDisk *
SVC remote copy consistency group	Consistency group of remote copy relations *
SVC cluster partnership	Copy- partnership linking of two SVC clusters *

* For more details about the object type please see: "IBM SVC software installation and configuration guide

1.10.21 SVC virtualization

Units of SVC/V7000 virtualization layer.

SVC extent	Storage area which connects the SVC MDisks with SVC VDisk copies within a SVC MDisk group *
SVC node	SVC nodes or V7000 controller nodes.*

SVC IO group	Group of two SVC/V7000/V3700 nodes, which together provide the IO of a SVC VDisk´s subset *
SVC cluster	Groups all SVC objects belonging to one cluster. More than one SVC cluster can be managed by BVQ. *

* For more details on the object type: please see "IBM SVC software installation and configuration guide".

1.10.21.1 SVC node

SVC nodes or V7000 controller nodes.

Groupes:

- Indirect: via the relation "preferred node" of a SVC VDisk

Is grouped by: (configurable object relations)

- BVQ room
- SVC IO group

Configurable attributes

- A SVC node can be allocated to a BVQ room.

1.11 Object editors

Via the edit objects function, objects can be enriched with other information and linked with each other.

What are object editors?

Some SVC-object types can be extended with business- or application-attributes. Some can be allocated to groupings. For these actions BVQ GUI-dialogs for editing are used. In this way, e.g. volumes in organization-units of a company or storage-classes in BVQ can be displayed.

Examples: Volumes can be mapped in organizational units of a company, or store classes can be displayed in BVQ.

How are the editors called?

1. Via the menu "object-lists" the tables views can be reached fot all editabe object types. Via right click on the row of an object the context menu is opened and leads via "edit" to the appropriate editor.
2. Via the context menu of a treemap-object
3. Via the context features in the list of related objects in the object editors.
4. Via the context menu of the option-dialog in the performance view.

Can several objects be edited at the same time?

If multiple objects are selected in tables view or in the treemap at the same time (hold key [CTRL] when selecting) and the option "edit" is used afterwards, the attributes or links of all objects can be changed in one draft.

What editors are available?
An overview of all BVQ object types and their order can be found in object types.
Object editors are only available for the following object types with BVQ-relevant attributes:

Tabelle 76 Editable object types

Name	Description
BVQ locations	
BVQ room	Datacenter room within a site, which contains servers, storage subsystems and nodes.
BVQ site	Datacenter site containing one or more datacenter-operating rooms.
Memory utilities	
BVQ disk type	Disk type found in disk arrays of storage subsystems.
BVQ disk array	Array (RAID) of one disk type in a storage subsystem.
BVQ controller	Represents a storage subsystem - cluster crossing group of SVC controllers.
SVC controller	Represents sub-units (port cards / controller-nodes with different WWNN) of a storage subsystem.
SVC MDisk	Corresponds to a logical disk of a storage subsystem (LUNs), which is assigned to a SVC/V7000
SVC MDisk group	Storage pool composed of SVC MDisk (logical disks)
SVC virtualization	
SVC node	SVC nodes or V7000 controller nodes.

Name	Description
Memory consumer	
SVC VDisk	SVC/V7000 volume (virtual disk).
SVC VDisk copy	SVC VDisk copy corresponds to SVC object type VDisk copy.
SVC host	Server which is connected to a SVC/7000 cluster and can be assigned to the SVC VDisks.
BVQ organizational units	
BVQ VDisk group	Volume group - groups SVC VDisk of the same application and quality level.
BVQ application	Application (e.g. Oracle, Windows File Service, Oracle DB Service, SAP DB Service) grouping BVQ VDisk groups.
BVQ application group	Application group - groups applications e.g. for the representation of specialized procedures, business processes etc.
BVQ cost center	Cost center - combines application groups e.g. cost carriers, customers, and departments of a company.
BVQ placement rules	
BVQ mobility zone	Group of storage pools to which a sub-set of assigned volumes can be placed.
BVQ VDisk group relation	Specifies a rule for the placement of two VDisk groups having the same usage context.
BVQ service level management	
BVQ storage class	Storage class - group of storage pools with the same quality classification.
BVQ SLA	Service level agreement - associates a BVQ VDisk group with a BVQ data class and details on prices (€ / GiB) and run time.
BVQ data class	Data class - arranges volume-groups with equal quality requirements
BVQ DC to SC placement rating	Requirements for placement combinations of data- and storage class

1.11.1 Object-editors

Via the edit function, objects with further information can be enriched and linked among each other.

What purpose do object-editors serve?

Some SVC object types can be extended with business- or application-attributes, or be allocated to groupings. For these actions there are dialogs for editing the BVQ GUI.

In this way, e.g. volumes in organization-units of a company or storage-classes in BVQ can be displayed.

How are the editors called?

1. Via the menu "object-lists" you arrive at the tabular views for all editable object types. Via right-click on the line of an object, the context-menu opens, which leads to the suitable editor via "edit".
2. Via the context-menu of an object in the treemap.
3. Via the context-functions of the list of linked objects in the object-editors.
4. Via the context-menu of the option-dialog of the perfomance view.

Can multiple objects be edited simultaneously?

If multiple objects are selected simultaneously in the tabular view or treemap (keep button [CTRL] pressed at the selection) and "edit" is chosen, the attributes or connections of all objects can be changed in one draft.

Which editors do exist?

An overview about all BVQ object types and their arrangements can be found under object-types.

Object editors are only available for the following object types with BVQ-relevant attributes:

Name	Description
BVQ locations	
BVQ room	Datacenter-room within a location, in which servers, storage subsystems and nodes are located.
BVQ site	Datacenter-location, in which one or multiple datacenter-operating rooms are arranged.
Storage provider	
BVQ disk type	Disk type, which is included in the disk arrays of the storage subsystems.
BVQ disk array	Combination (RAID) of disks of one type in a storage subsystem
BVQ controller	Represents a storage subsystem - cluster-overlapping group of SVC

Name	Description
	controllers.
SVC controller	Represents part-units (port-cards / controller-nodes with different WWNN) of a storage subsystem.
SVC MDisk	Complies with a logical disk of a storage subsystem (LUN), which is allocated to a SVC/V7000.
SVC MDisk group	Storage pool consisting of SVC MDisks (logical Disks).
SVC virtualization	
SVC node	SVC accounts or V7000 controller node.
Storage utilizer	
SVC VDisk	SVC/V7000 volume (virtual disk).
SVC VDisk copy	SVC VDisk copy corresponds to SVC object type VDisk copy.
SVC host	Server, which is connected with a SVC/V7000 cluster and can be allocated to the SVC VDisks.
BVQ organization units	
BVQ VDisk group	Volume group - arranges SVC VDisks with the same application-quality level.
BVQ application	Application (e.g. Oracle, Windows File Service, Oracle DB Service, SAP DB Service) groups BVQ VDisk groups.
BVQ application group	Application group - arranges applications e.g. for the representation of specialized procedures, and business processes.
BVQ cost center	Cost unit - groups application groups e.g. benefactors, customers, departments of a company.
BVQ placement regulations	
BVQ mobility zone	Group of storage pools, on which a subset of allocated volumes can be placed.
BVQ VDisk group relation	Requirement of a regulation for the placement of two VDisk groups which are related in their application.

Name	Description
BVQ service level management	
BVQ storage class	Storage class - group of storage pools with the same quality classification.
BVQ SLA	Service level agreement - when a BVQ VDisk group is linked with a BVQ data class containing details about prices (€/GiB) and duration.
BVQ data class	Data class - groups volume-groups with the same quality-requirements.
BVQ DC to SC placement rating	Requirements for placement combinations of data- and storage class

Editable object types

1.12 Use cases

Practical applications in BVQ are shown in this chapter.

1.12.1 Qualification of the performance edition and the single extension-packages

In this chapter the qualification of the performance edition and the single extension-packages is described.

In the premium edition, several qualification-steps, which are described here, have to be carried out. Both extension-packages are part of these steps.

The qualification of the "MDisk to Array" is only possible in the entry edition. Please see for that: http://wiki.sva.de/#Qualifizierung-QualifizierungDiskArrays

Therefore, entry edition is not mentioned in this chapter.

1.12.1.1 Indication of unqualified objects

To obtain an overview about the current qualification-status the GUI treemaps can be used.

1. Placing object colors on BVQ qualified.
2. All treemap objects, which are not qualified yet, are colored red.
3. Select the desired object e.g. an object-combination in the treemap navigator.

4. Then with a right-click the unqualified object can be editted/qualified.

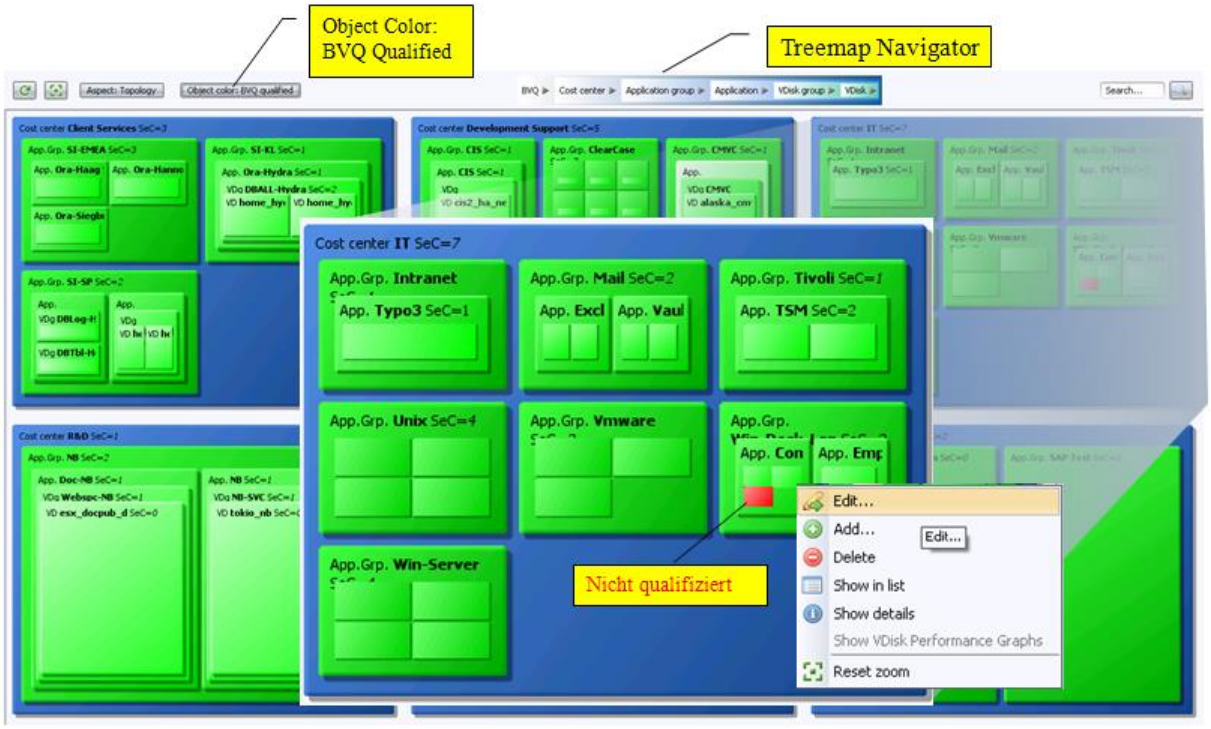

Use the **step-by-step procedure** in **favorites: qualify and edit objects for inspecting the qualification-progress** or to carry out the qualification in single treemaps. The sequence of the qualification might (if necessary) slightly deviate from what is told in the User Manual Guide. It is important that valid combinations of the BVQ exist.

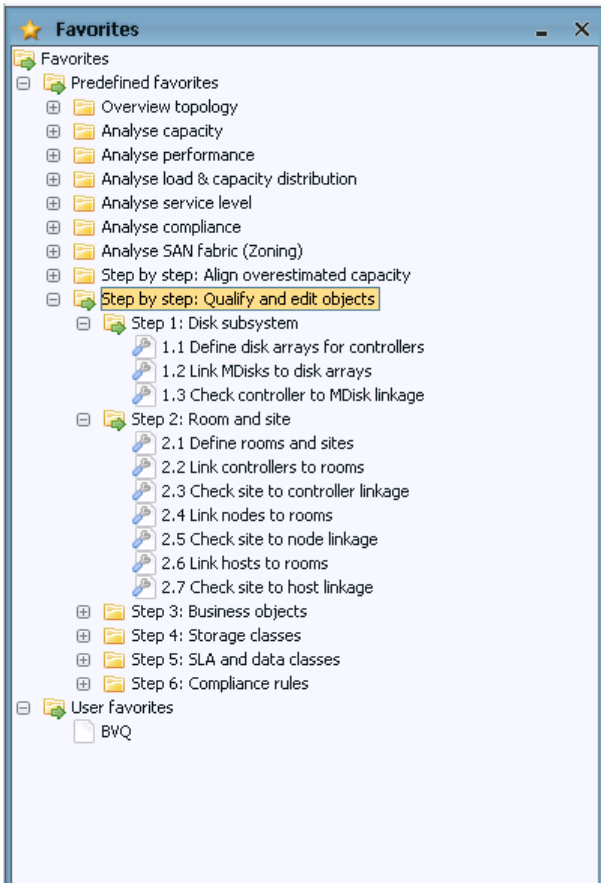
Favorites
Favorites
Predefined favorites
Overview topology
Analyse capacity
Analyse performance
Analyse load & capacity distribution
Analyse service level
Analyse compliance
Analyse SAN fabric (Zoning)
Step by step: Align overestimated capacity
Step by step: Qualify and edit objects
Step 1: Disk subsystem
1.1 Define disk arrays for controllers
1.2 Link MDisks to disk arrays
1.3 Check controller to MDisk linkage
Step 2: Room and site
2.1 Define rooms and sites
2.2 Link controllers to rooms
2.3 Check site to controller linkage
2.4 Link nodes to rooms
2.5 Check site to node linkage
2.6 Link hosts to rooms
2.7 Check site to host linkage
Step 3: Business objects
Step 4: Storage classes
Step 5: SLA and data classes
Step 6: Compliance rules
User favorites
BVQ

1.12.1.2 Qualifying the performance edition

In the first step of the qualification, the topology-objects are qualified. They comply with the whole qualification-range in the performance edition and are part of the qualification-range of the premium edition.

For the first step at the qualification, the following information is needed.

- **Sites with rooms**
  From all hosts and storage systems which have to be qualified, the site and room information has to be available.
  Where are the storage systems, SVC nodes and hosts located?
  (location and room)
- **Storage system information**
  Evaluation (lowrange, midrange, highend), cachesize, MTM information (specific technical description) and serial number.
- **Disk type**
  Which disk types are applied (SATA, SAS,SSD, 10kRPM, 15kRPM, max. IO\s)?

- **Disk array Information**
  Reading out all existing disk arrays of the storage systems: raid level, disk type, amount of disks, controller, MDisks
- **Determining controller**
  Cache size and classifying it as enterprise, midrange or lowrange.
- **Server / hosts**
  Location and contact
- **SVC nodes:**
  Location

Hint: discovery of technical Information

- Example: DS3500
  - Opening DS storage manager
  - Calling up profile
  - Reading out Information from sections (controller cache, MTM, serial numbers, raid information)

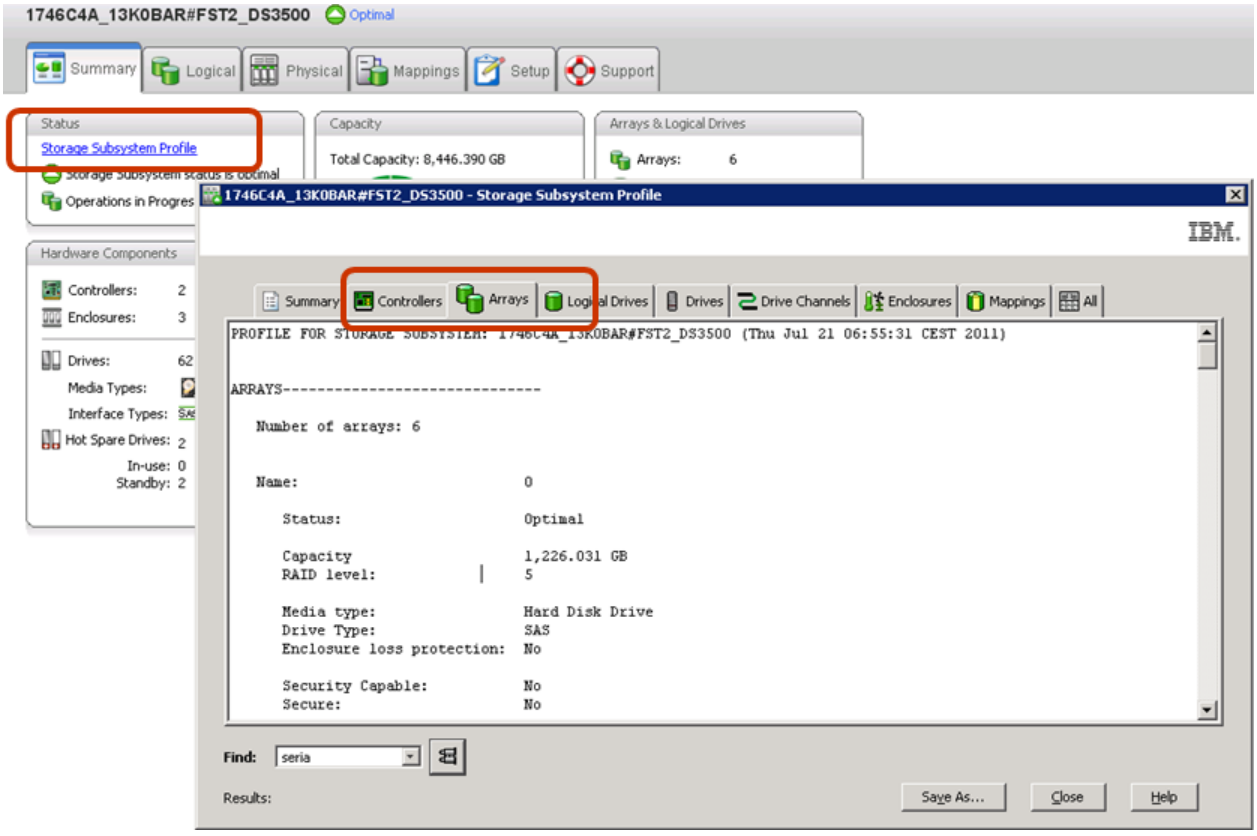

Qualifying site & room

1. Start the qualification in the task menu with edit -> site and click on add.

2. Enter a name for the site. Repeat this procedure with every site.

3. Carry out the same steps for the site room (**edit->room**) and allocate to each room the appropriate site.

Clue: the comment-field can be utilized for additional Information, like e.g. personal contacts.

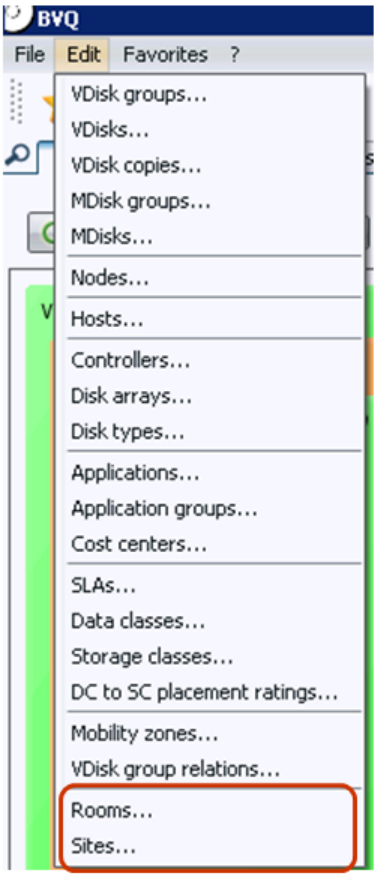

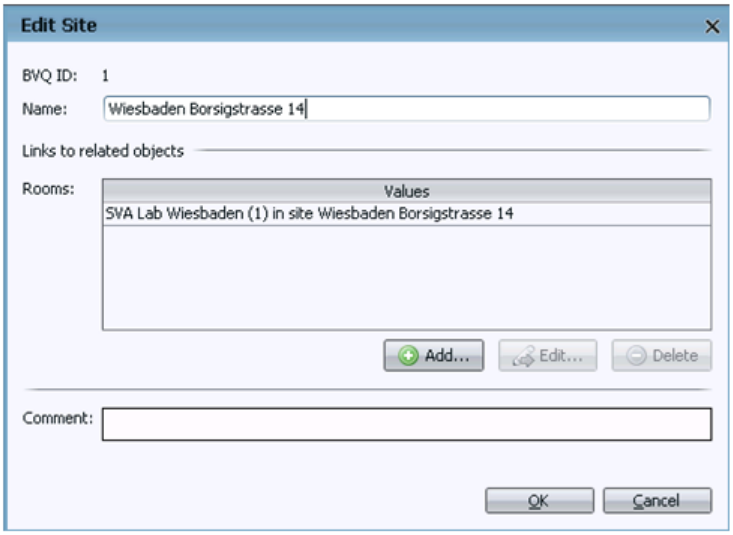

Qualifying disk types

For the qualification of disk types, the disk-attributes have to be determined first.

- **Class:** concerning lowrange (ATA, SATA, PATA), enterprise (FC, SAS, SCSI) or SSD storage ?
- **Capacity:** which capacity do the disks have?
- **Rotation Speed:** which configuration do the disks have?
- **Max est. I/O rate:** Is the speed determined by rotation? This value can also be adapted.

Start the qualificaton in the task menu with edit -> disk types and click on add.

Qualifying disk controller

For the qualification of the controller you select the correspondent class and room and register the cache Size. The room Information must be entered in a way so that the object (controller) can be fully qualified. The cache size, MTM and serial number have to be gathered by the controller. Determining controller cache size and classifying it as enterprise, midrange or lowrange.

Start the qualification in the task menu with edit -> controllers... and select the controller which has to be qualified.

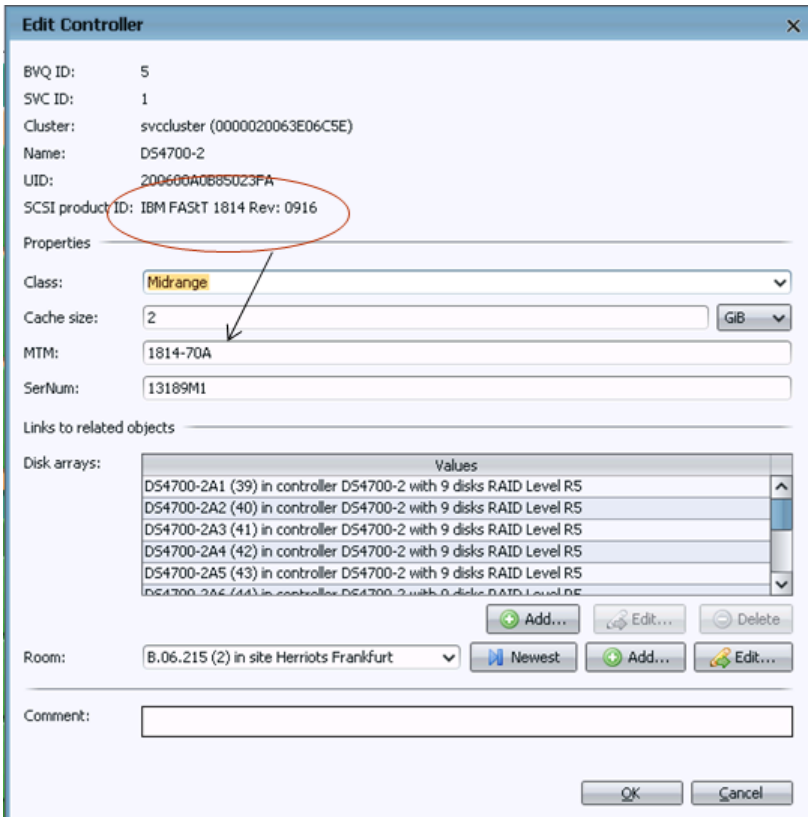

Qualifying disk arrays

In the disk arrays the performance attributes of the hardware (MDisks) which lie underneath, are gathered. The gathered data of the disk arrays can later be used to generate reports (e.g. table of all raid 5 arrays).

At the Storwize V7000, disk arrays are automatically created for each MDisk. No user-action is necessary. If multiple MDisks have to be summarized into one disk array, the remaining disk array can be discarded and new ones can be created.

You enter disk array name, raid level, disk type, number of disks (without hotspares) and the controller.
The estimated I/O density is calculated by the specifications of the raid level and disk type (max est. I/O rate). In the case of deviation this value can be overwritten.

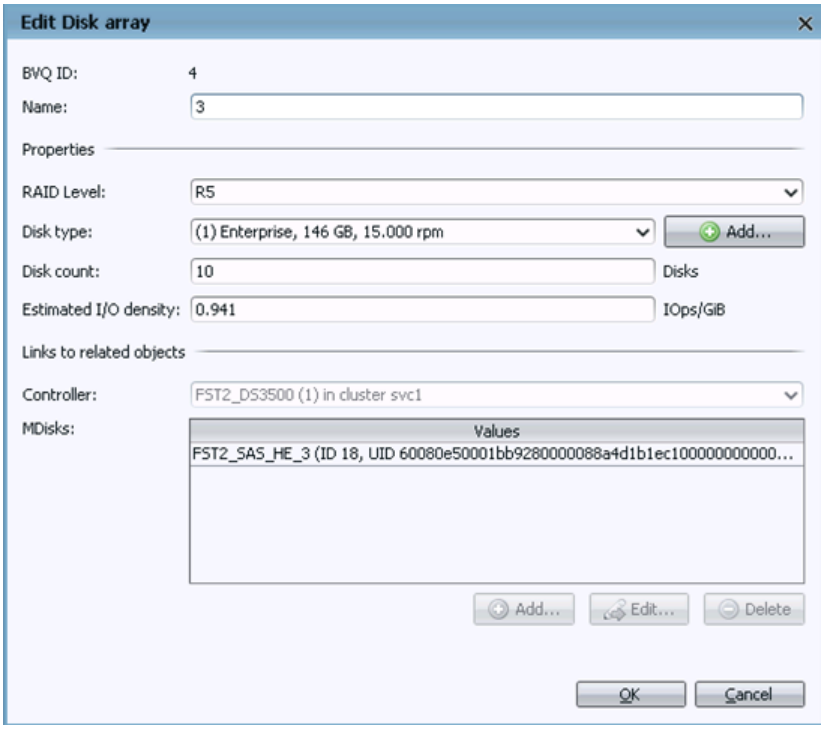

Utilize the treemap path stored in favorites for more efficient visualization of controllers, disk array and mDisk:
Favorites-> predefined favorites -> step by step: qualify and edit objects -> define disk array for controllers.

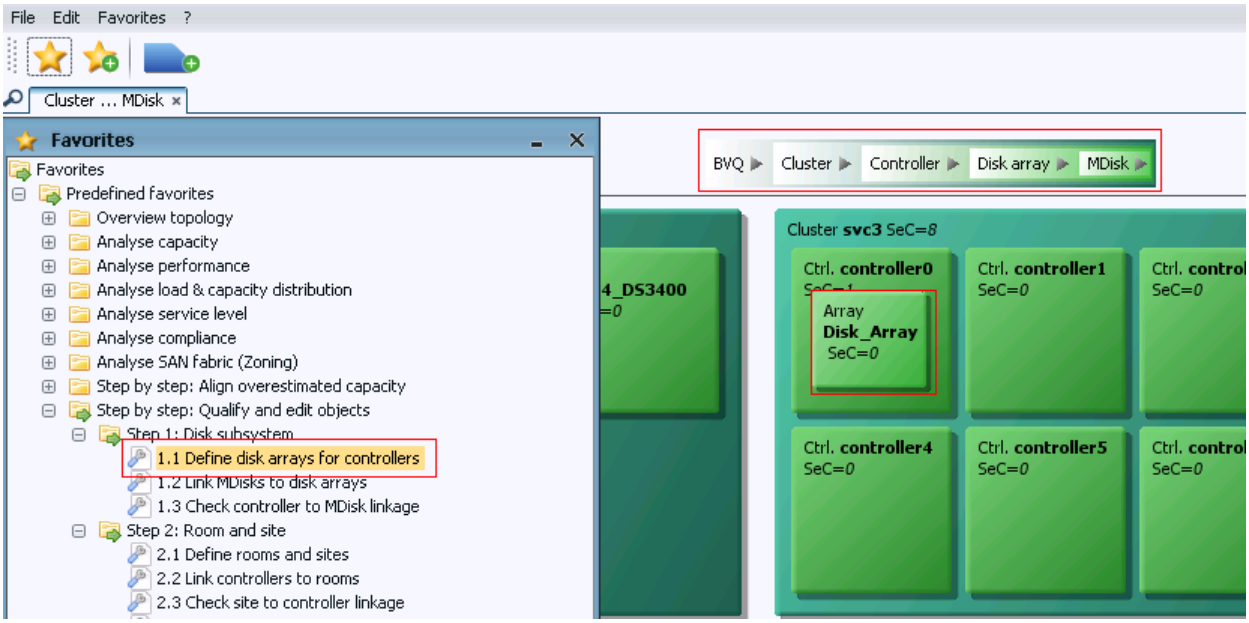

Switch to the treemap Path subsequently: BVQ -> cluster -> controller -> MDisk -> disk array and switch to object color: BVQ qualified. MDisks which are not allocated yet to any array are displayed in red.
In the treemap a new disk array can directly be created for the controller and subsequently the MDisks can be allocated to the new array by editting the controller.

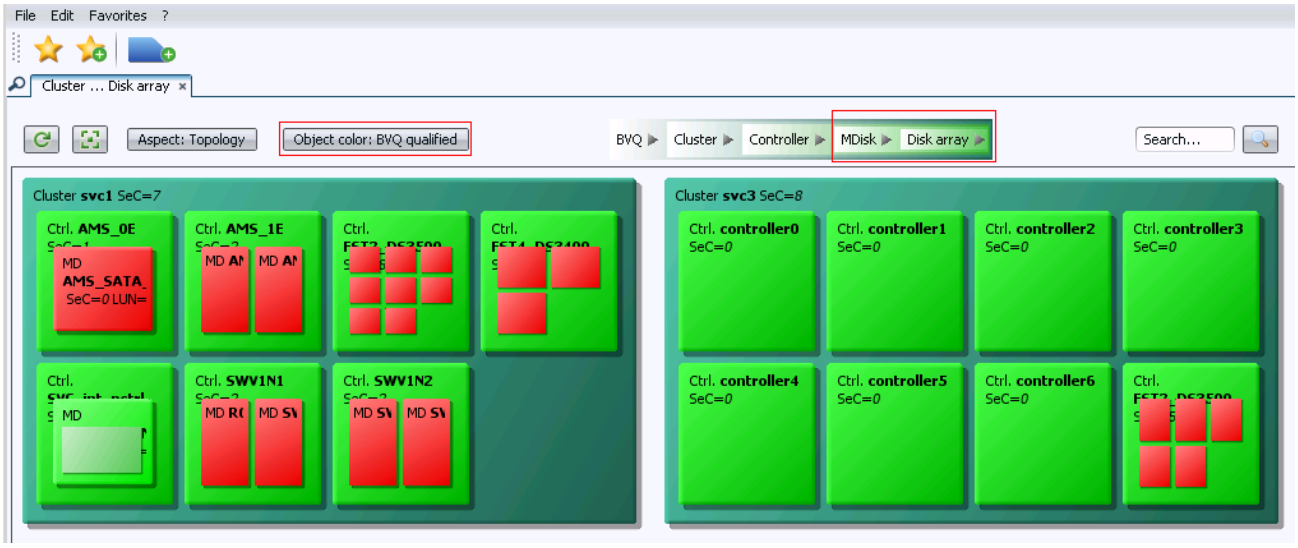

Qualifying hosts

The hosts are gathered via the BVQ scanning-procedure. The specification of the room is required for the qualification. Furthermore the comment field can be applied for contact (already mentioned above).

The qualification is started via the task menu with *edit -> hosts* and selecting the host that has to be qualified.

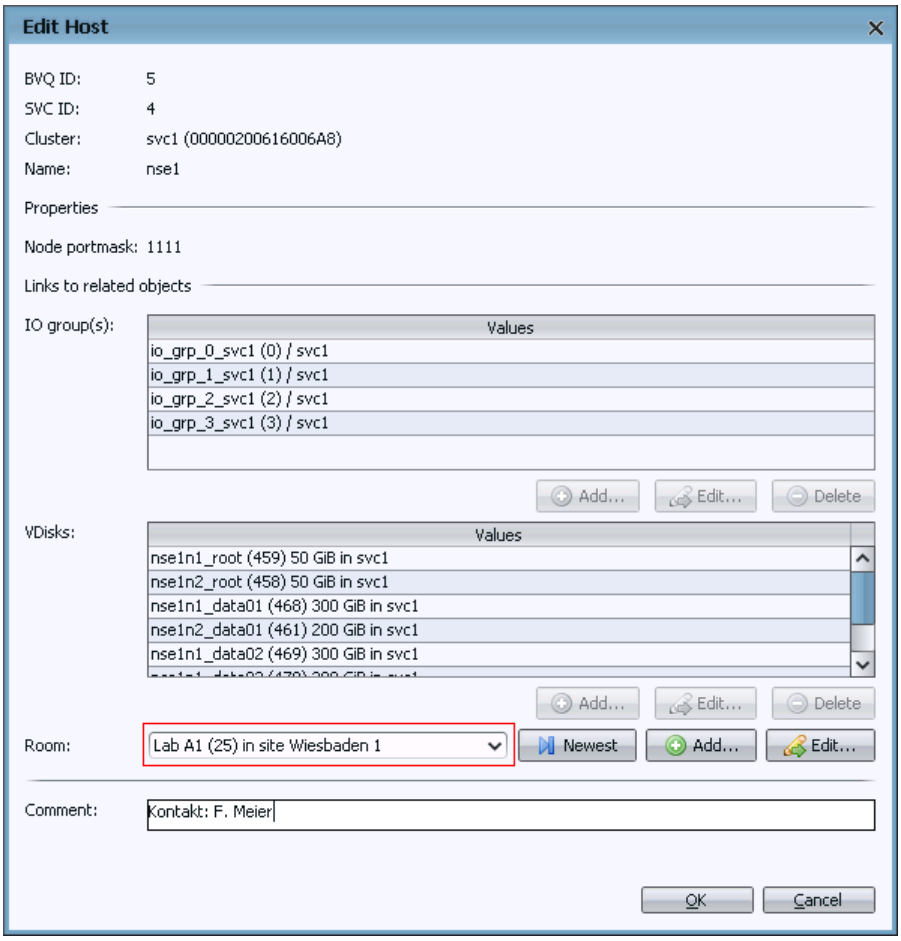

Qualifiying SVC nodes

The nodes are also gathered via the BVQ scanning-procedure. The specification of the room is required for the qualification.

The qualification is started via the task menu with *edit -> nodes* and selecting the node which has to be qualified.

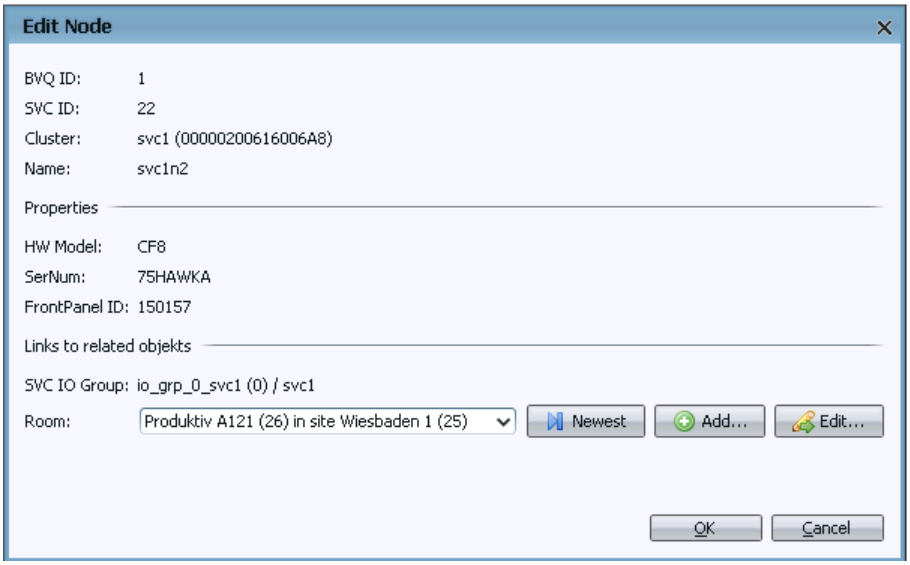

1.12.1.3 Qualifying BVQ accounting package

The accounting package is part of the premium edition and is optionally available for the performance edition. The qualification-steps which are described in this collumn can be executed, if the premium edition or the performance edition with the accounting package is used.

For the accounting package information about the company is needed, and thereby the following questions are asked: "Which applications are implemented in the company-environment?" or "How can these applications be sensibly gathered in an "invoice-technical" way?
For this reason it is necessary in BVQ to summarise and expediently link the correspondent information subsequently.

- What are my cost units?
  E.g. cost unit: human resources department, customer xy or software development
  -> This information is generally registered in the group "cost center".

- Which application groups operate in single cost units?
  E.g. salary accounting, centralized chronology
  -> This information is generally registered in the "application group".

- Which applications operate in the application groups?
  E.g. SAP Prod, SAP Test, notes, Citrix or the development-software
  -> This Information is generally registered in the group "application".

- Are the applications distributed among different systems?
  E.g. tablespaces, backup
  -> This Information is generally registered in the group "VDisk group".

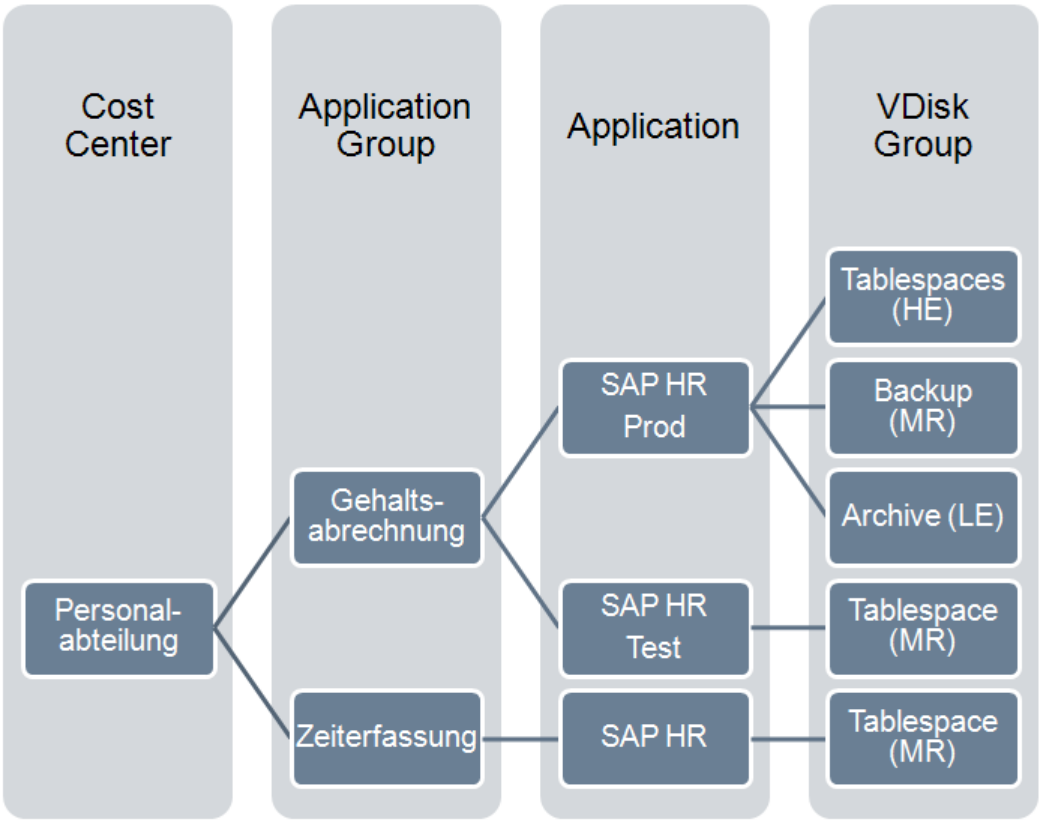

The VDisk groups in this picture show the basis of the accounting package. With the VDisk groups the SLA's, data-classes, mobility zones, VDisks and the VDisk group relations are connected.

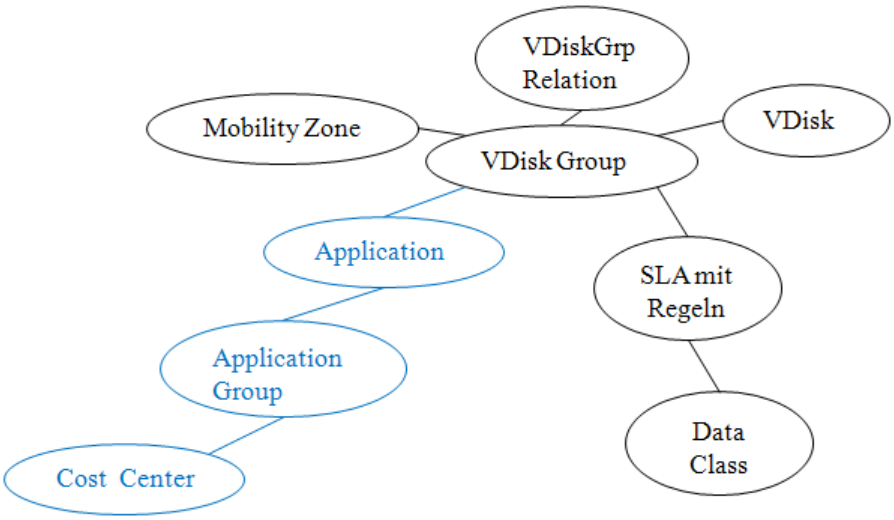

Qualifying cost center, application group, applications, VDisk group, VDisk

The qualification is for instance meaningfully arranged. It starts with the cost center and continues with the application groups, applications and finally carries out the qualification of the VDisk groups. Thereby a new object can be connected to an already existing object.

Object cost centers

The qualification is started via the task menu with *edit -> cost centers...* and with a click on add to create different cost centers.

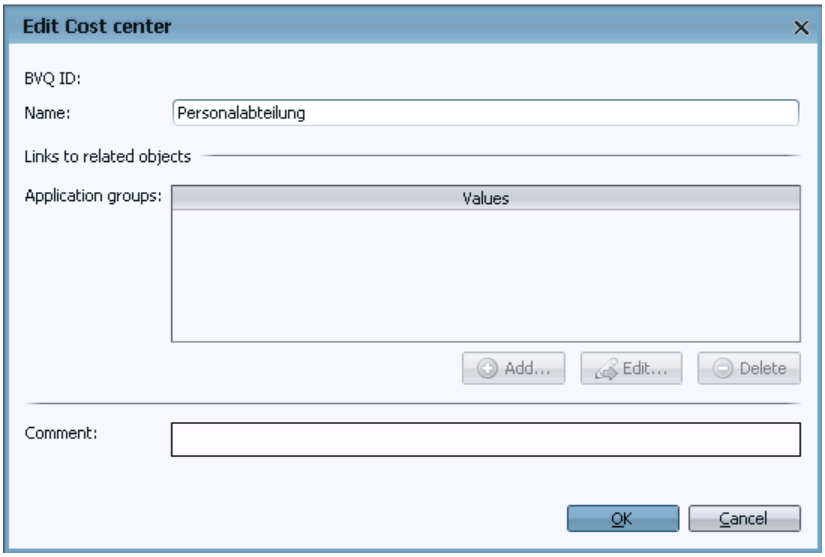

Object application groups

The qualification is started via the task menu with *edit -> application groups...* and with a click on add to create the application groups. Allocate the suitable cost center to each application group.

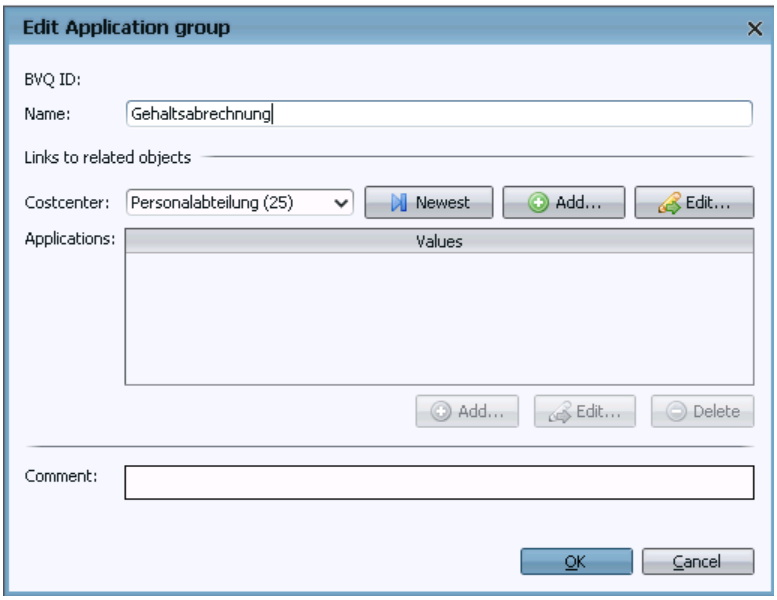

Object application

The qualification is started via the task menu with *edit -> applications...* and with a click on add to create the applications. Allocate the suitable application group to each application.

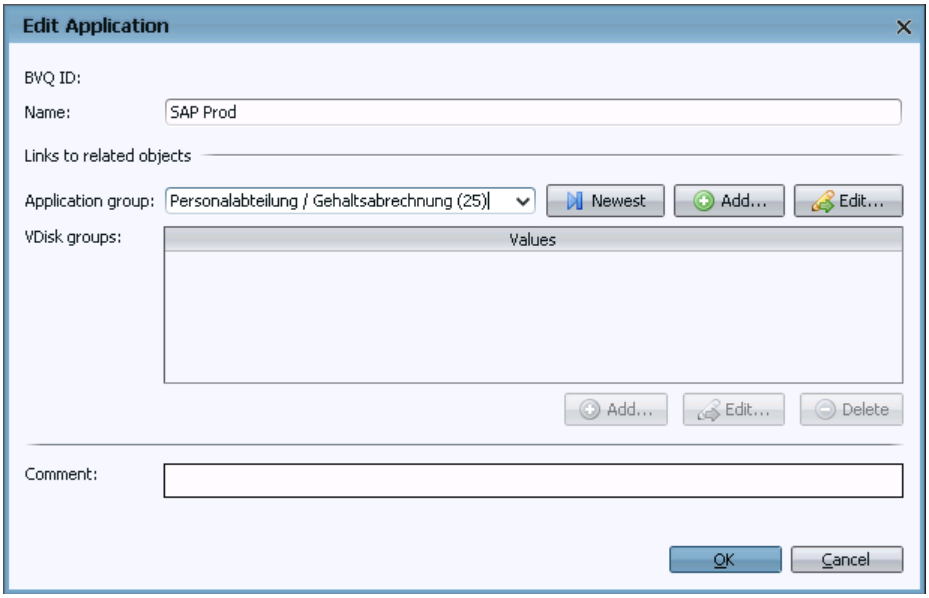

Object VDisk groups

The qualification is started via the task menu with *edit -> VDisk groups...* and with a click on add to create the VDisk groups. Allocate the suitable application to each VDisk group.

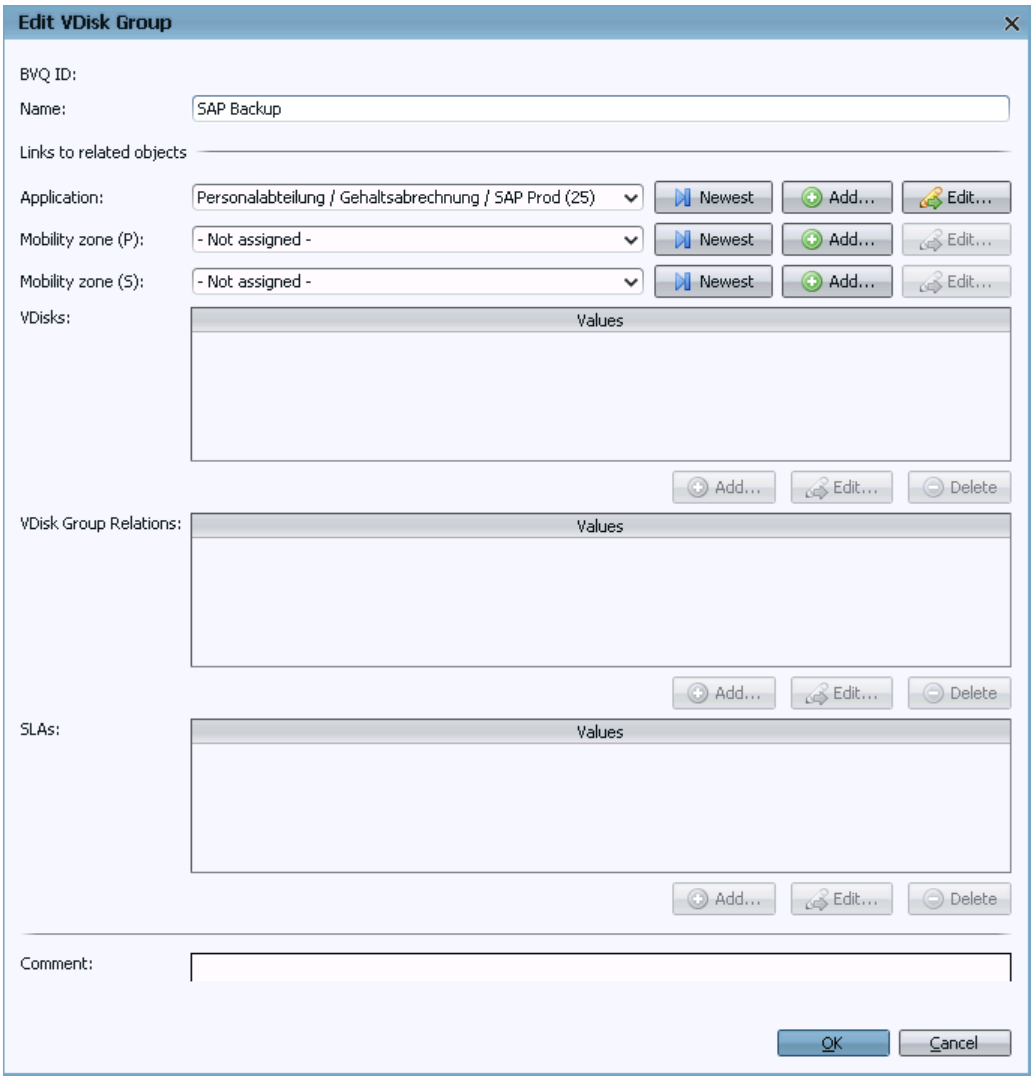

VDisk

The qualification is started via the task menu with *edit -> VDisks...* and selecting a VDisk which can be connected with a previously created and suitable VDisk Group.

Utilize the name filter to locate VDisks faster.

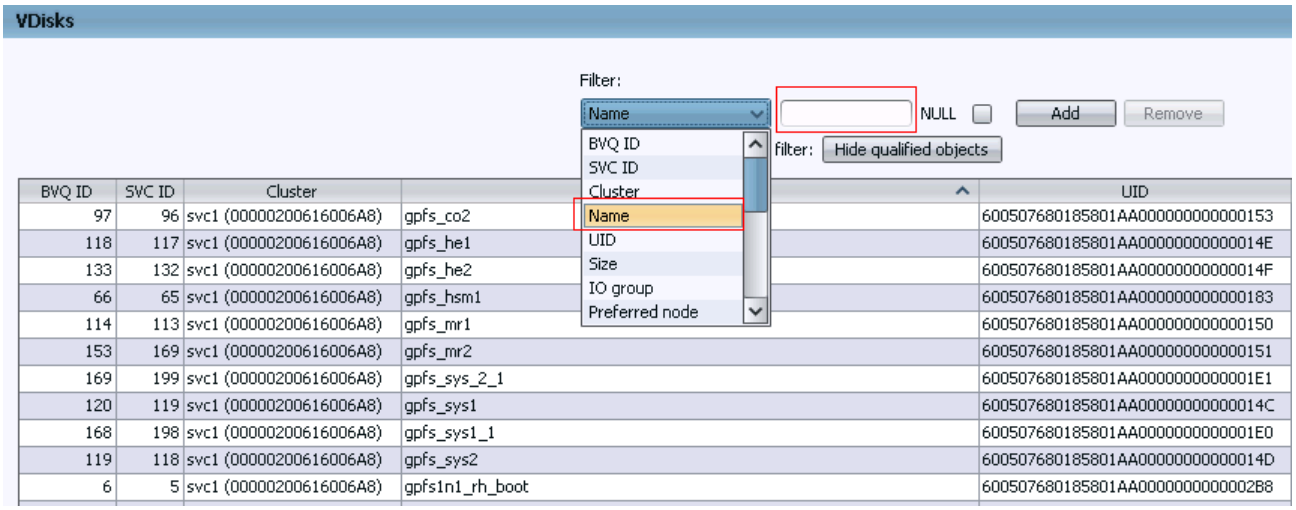

Then the valid VDisk group is allocated to it. The qualification is carried out for the other VDisks as well.

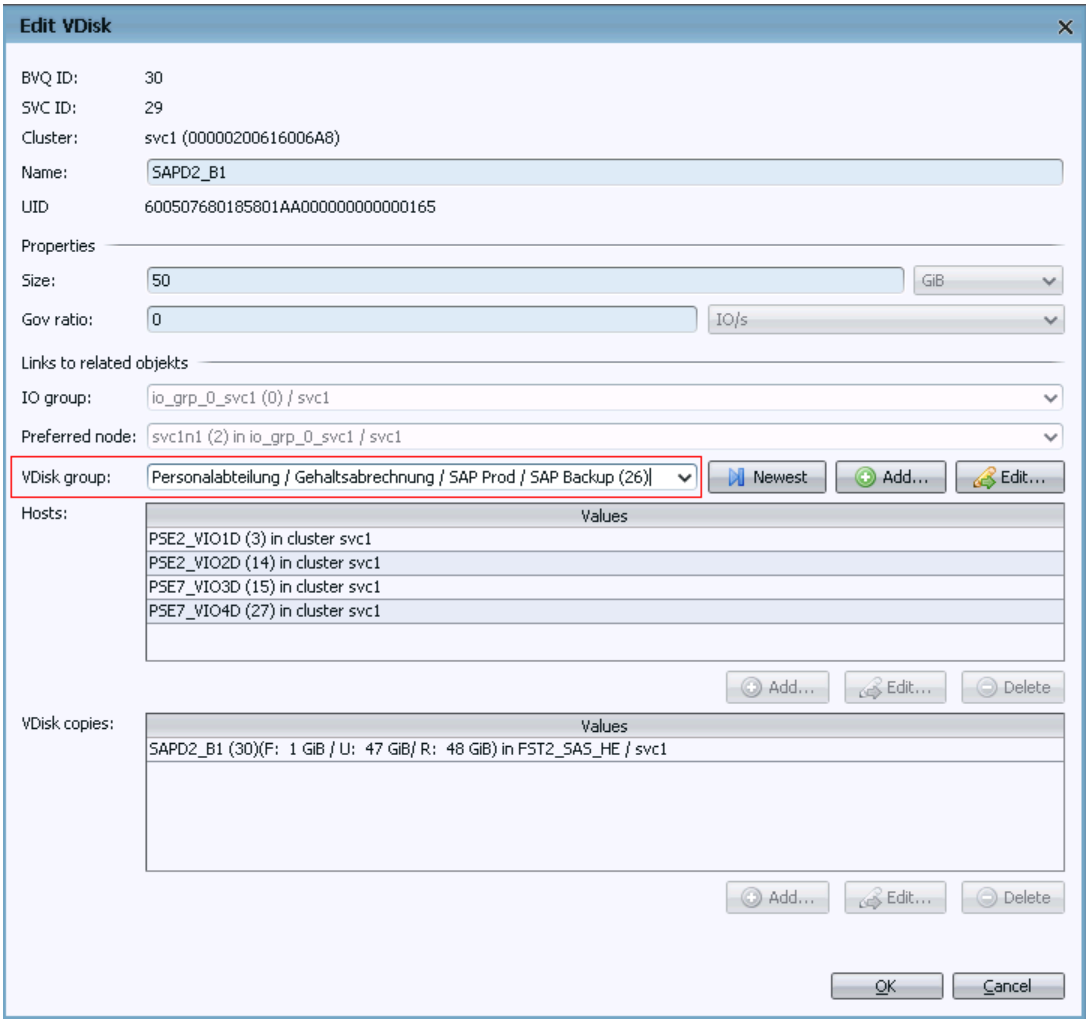

After these executed qualification steps, the first evaluations in the BVQ can take place (e.g. the capacity-view / performance view on a cost center).

Capacity view:

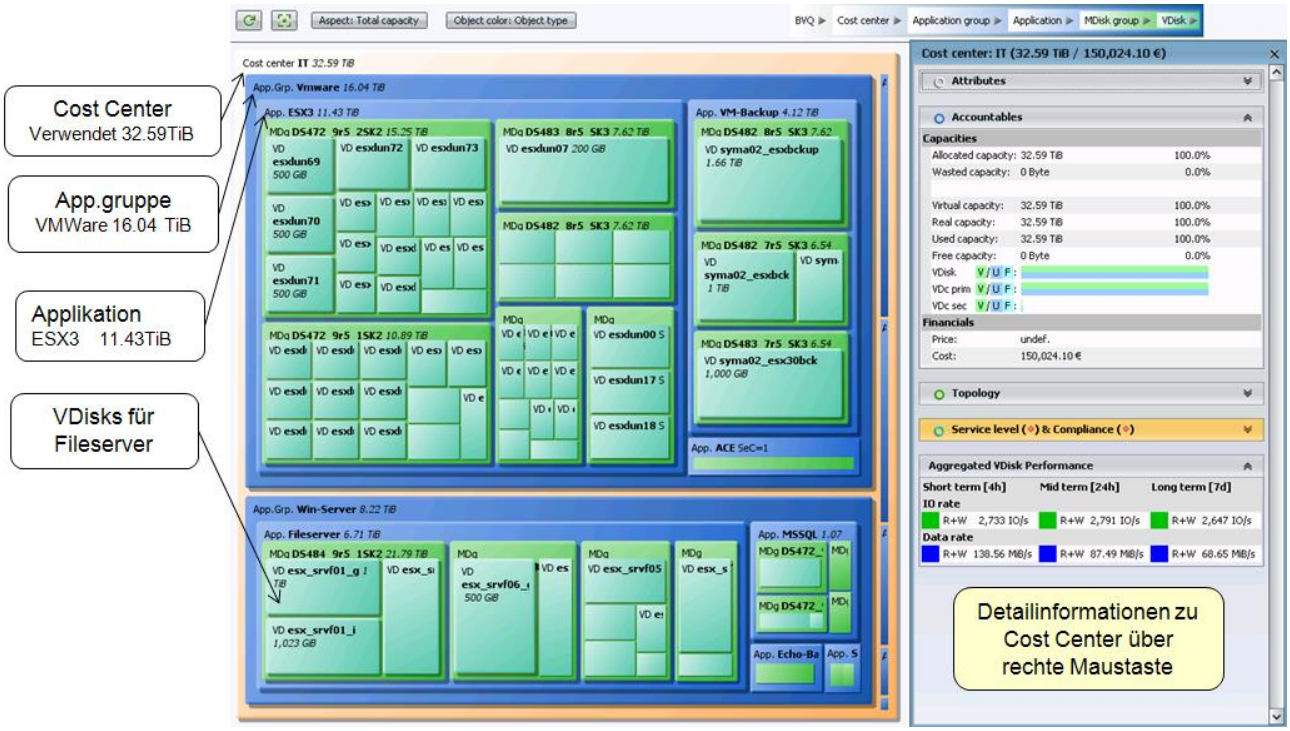

Performance view:

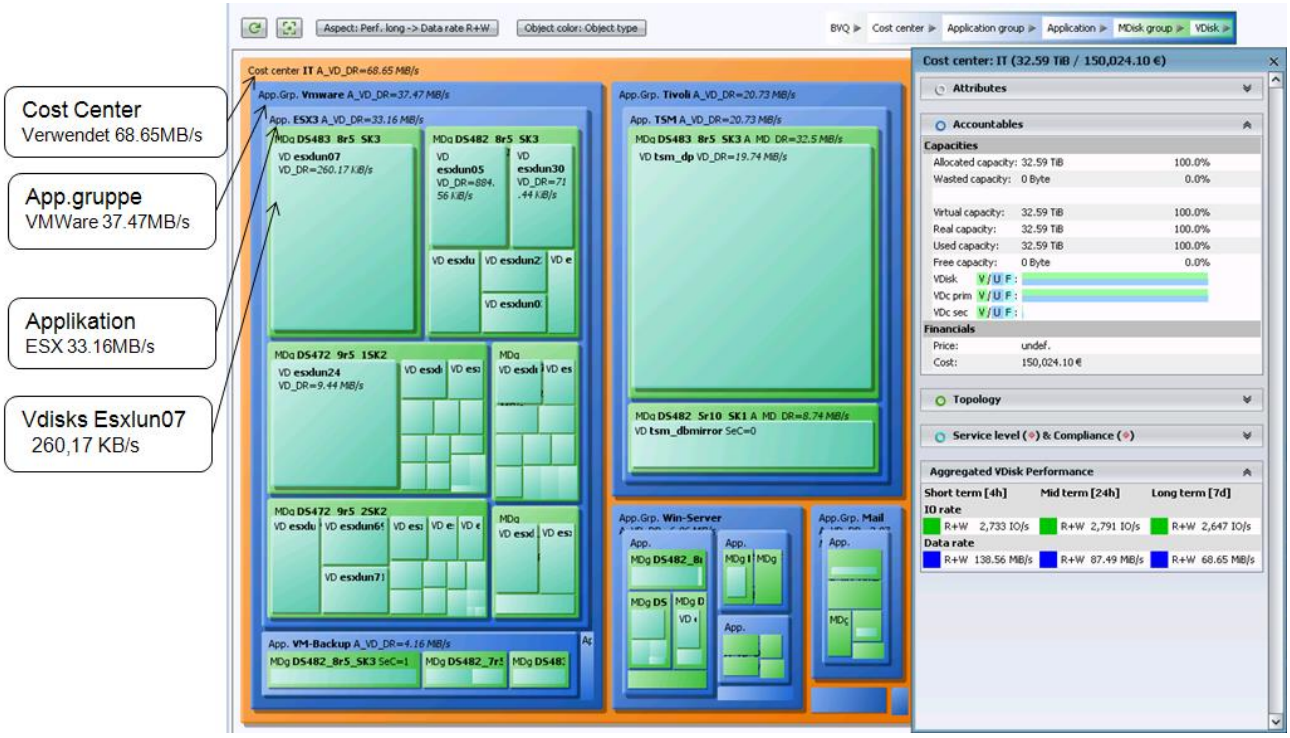

Qualifying mobility zones

You now start Beginnen Sie nun die Mobility Zones einzurichten. Mobility Zones legen fest, auf welchen Storage Pools die VDisks einer VDisk Gruppe platziert werden sollten. Somit können Regeln für Ablageorte von VDisks festgelegt werden, die dann beispielsweise auf Ihre Einhaltung hin überprüft werden können (VDisk X sollte nur auf den MDisk groups Y & Z abgelegt werden).

Erstellen und MDisk Gruppen zuweisen
Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Mobility zones...**und klicken Sie auf *Add**.

Mobility Zones den VDisk Gruppen zuweisen
Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> VDisk groups...* und wählen sie eine passende VDisk Gruppe aus.
Eine Mobility Zone kann hierbei für mehrere VDisk Gruppen verwendet werden.
Geben Sie für die Mobility Zone (P) die Mobilty Zone an, in welcher sich die primären Kopien der VDisks in der VDisk Gruppe befinden sollen.
Geben Sie für die Mobility Zone (S) die Mobilty Zone an, in welcher sich die sekundären Kopien der VDisks in der VDisk Gruppe befinden sollen. Somit können gespiegelte VDisks (VDisk Mirror) unterschiedlichen Mobility Zones zugewiesen werden.

Mobility Zone Compliance Check

Über **Object Colors: Mobility zone compliance** werden in der Einfärbung die falsch platzierten Objekte rot angezeigt. Korrekt platzierte Objekte werden immer grün angezeigt. Blaue Objekte weisen darauf hin, dass keine Complianceregeln bestehen.

Nutzen sie auch hier die Step by Step -> Compliance rules Funktion um schnell einen Überblick über die Compliance zu erhalten.

In der Premium Edition müssen sämtliche hier beschriebenen Qualifizierungsschritte durchgeführt werden, da beide Erweiterungspakete Bestandteil dieser sind.

Bei der Entry Edition ist einzig die Qualifizierung der "MDisk zu Array" möglich. Siehe dazu: #Qualifizierung Disk Arrays.
Auf die Entry Edition wird daher in diesem Kapitel nicht weiter eingegangen.

1.12.1.4 Display of unqualified objects

Einen Überblick über den aktuellen Qualifizierungsstand erhält man durch die Treemaps der GUI.

1. Object Colors auf BVQ qualified stellen
2. alle Treemap Objekte, die noch nicht qualifiziert sind werden jetzt rot eingefärbt.
3. Im Treemap Navigator das gewünschte Objekt bzw. eine Objektkombination auswählen
4. Das unqualifizierte Objekt kann mit einem Rechtsklick anschließend editiert/qualifiziert werden.

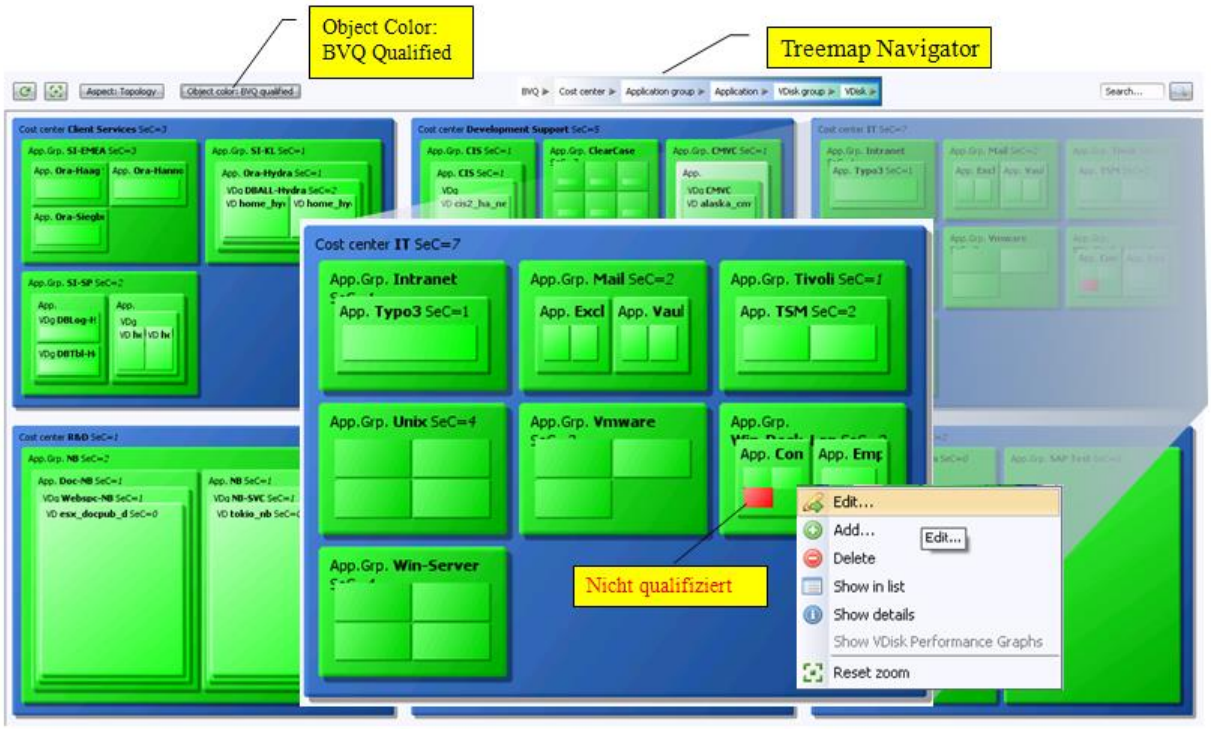

Nutzen Sie das unter **Favoriten** gespeicherte **Step by step: Qualify and edit objects** Verfahren um den Qualifizierungsfortschritt zu überprüfen oder direkt in den einzelnen Treemaps die

Qualifizierung durchzuführen. Die Reihenfolge der Qualifizierung kann ggf. von der im Benutzerhandbuch abweichen. Wichtig ist, dass am Ende gültige Kombinationen der BVQ Objekte bestehen.

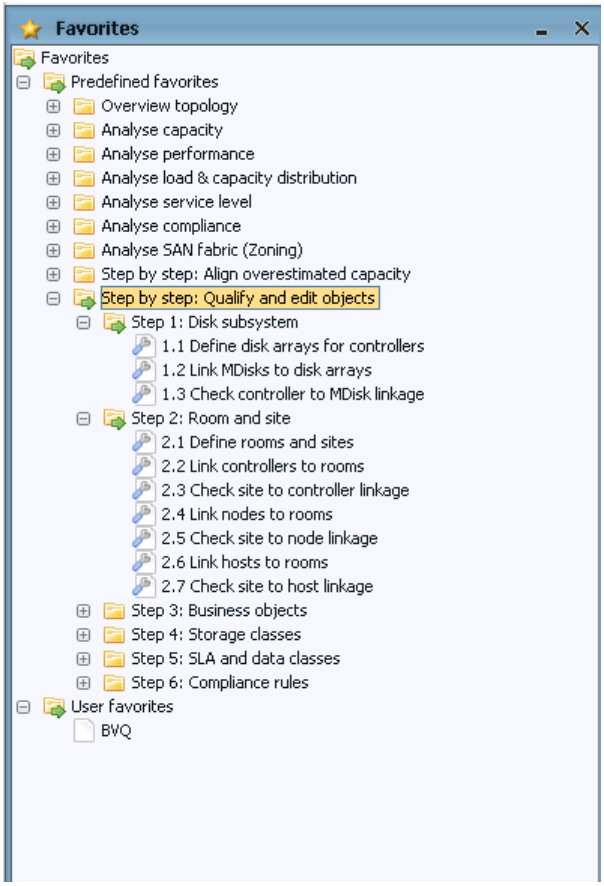
Favorites
Favorites
Predefined favorites
Overview topology
Analyse capacity
Analyse performance
Analyse load & capacity distribution
Analyse service level
Analyse compliance
Analyse SAN fabric (Zoning)
Step by step: Align overestimated capacity
Step by step: Qualify and edit objects
Step 1: Disk subsystem
1.1 Define disk arrays for controllers
1.2 Link MDisks to disk arrays
1.3 Check controller to MDisk linkage
Step 2: Room and site
2.1 Define rooms and sites
2.2 Link controllers to rooms
2.3 Check site to controller linkage
2.4 Link nodes to rooms
2.5 Check site to node linkage
2.6 Link hosts to rooms
2.7 Check site to host linkage
Step 3: Business objects
Step 4: Storage classes
Step 5: SLA and data classes
Step 6: Compliance rules
User favorites
BVQ

1.12.1.5 Qualifizierung Performance Edition

Im ersten Schritt der Qualifizierung werden die Topologie-Objekte qualifiziert. Diese entsprechen dem kompletten Qualifizierungsumfang in der Performance Edition und sind Teil des Qualifizierungsumfangs der Premium Edition.

Für den ersten Schritt bei der Qualifizierung werden die folgenden Informationen benötigt.

- **Sites mit Rooms**
  Von allen Hosts und Speichersystemen, die qualifiziert werden sollen, müssen die Site und Room Informationen vorliegen.
  Wo stehen die Speichersysteme, SVC Knoten und Hosts?
  (Ort und Raum)
- **Speichersysteminformationen**
  Einschätzung (Lowrange, Midrange, Highend), Cachesize, MTM Informationen (korrekte technische Bezeichnung) und Seriennummer.
- **Disk Type**
  Welche Disktypen werden verwendet (SATA, SAS,SSD, 10kRPM, 15kRPM, max. IO\s)?

- **Disk Array Informationen**
  Alle bestehenden Disk Arrays der Speichersysteme auslesen: Raidlevel, Disktyp, Anzahl Disks, Controller, MDisks
- **Controller** Cache Size ermitteln und als Enterprise, Midrange oder Lowrange klassifizieren.
- **Server / Hosts:** Standort und Kontakt
- **SVC Nodes:** Standort

Tipp: Auffinden technischer Informationen

- Beispiel DS3500
  - DS Storage Manager öffnen
  - Profile aufrufen
  - Informationen aus Sektionen herauslesen (Controller Cache, MTM, Seriennummern, Raidinformationen)

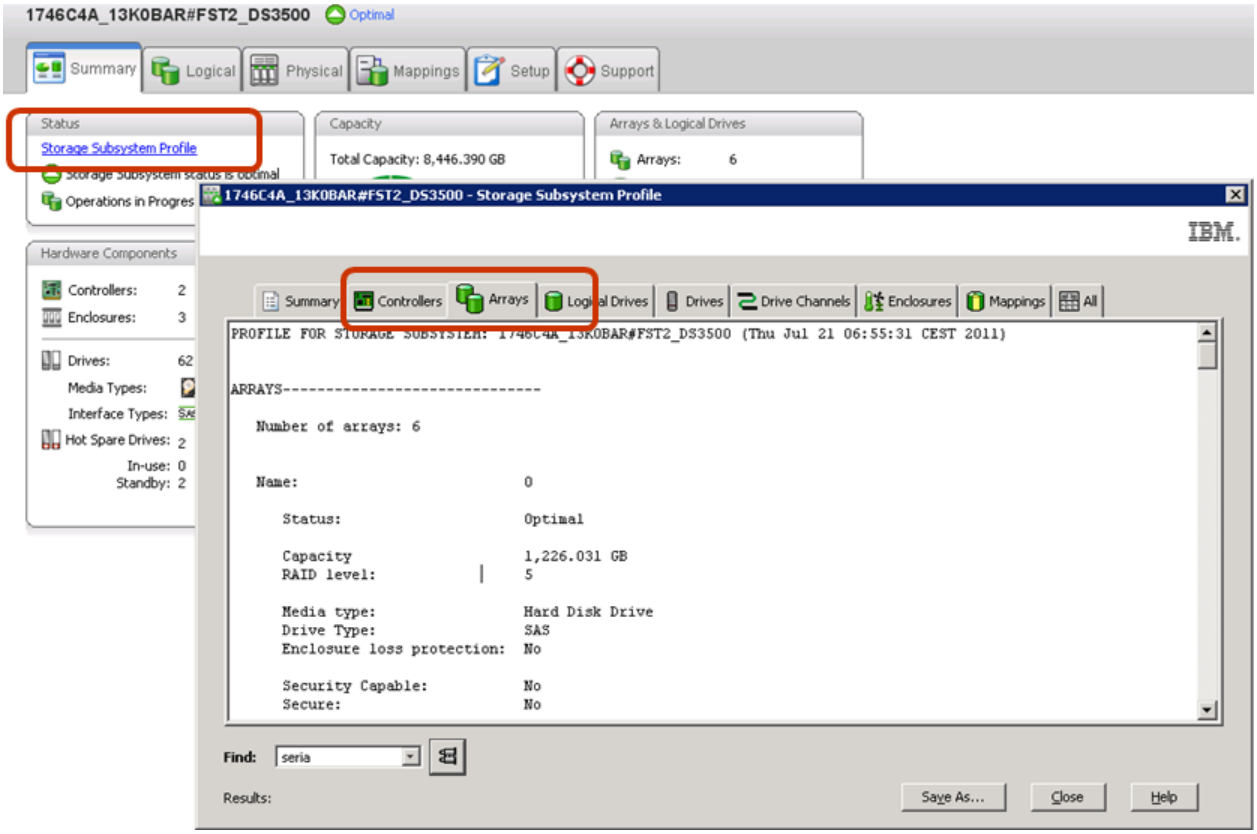

Qualifizierung Site & Room

1. Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Site*und klicken Sie auf *Add*.

2. Geben Sie einen Namen für die Site ein. Wiederholen Sie diesen Vorgang für jede Site.

3. Führen Sie dieselben Schritte für den Room aus (**Edit->Room**) und weisen sie jedem Room die dafür vorgesehene Site zu.

Hinweis: Das Comment-Feld kann für zusätzliche Informationen genutzt werden, wie z.B. Kontaktpersonen.

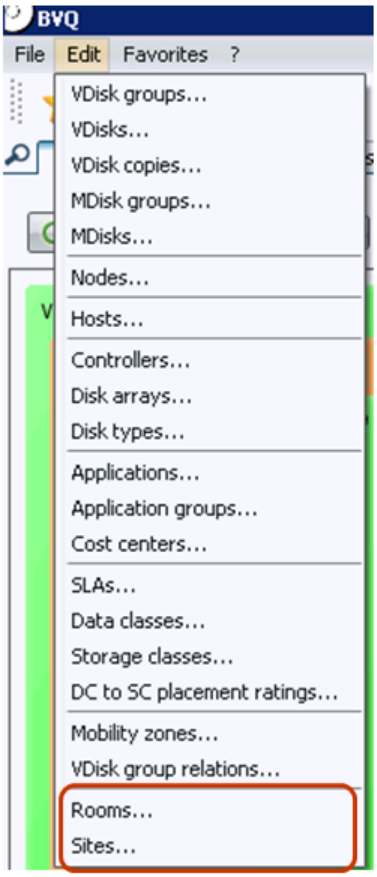

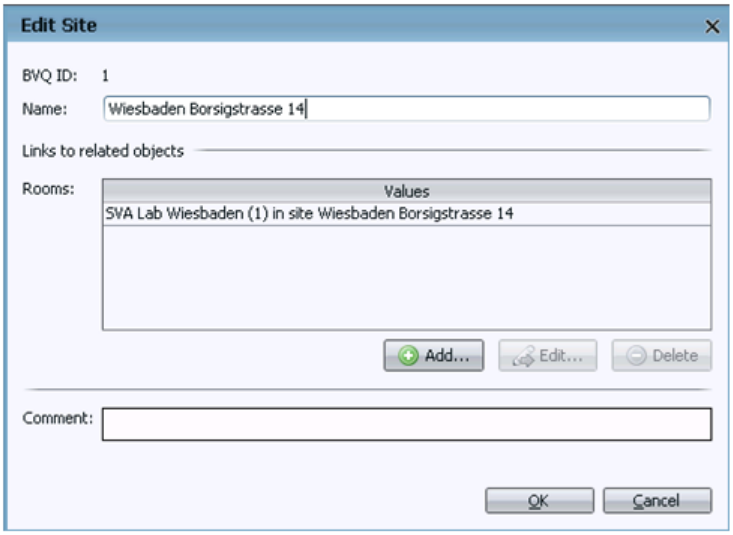

Qualifizierung Disk Typen

Für die Qualifizierung der Disk Typen müssen im Vorfeld die Disk-Eigenschaften erfasst werden.

- **Class:** Handelt es sich um Lowrange (ATA, SATA, PATA), Enterprise (FC, SAS, SCSI) oder SSD Storage ?
- **Capacity:** Über welche Kapazität verfügen die Disks ?
- **Rotation Speed:** Welche Umdrehungszahlen haben die Disks ?
- **Max est. I/O rate:** Wird aus Rotation Speed ermittelt. Dieser Wert kann auch angepasst werden.

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Disk types*und klicken Sie auf *Add*.

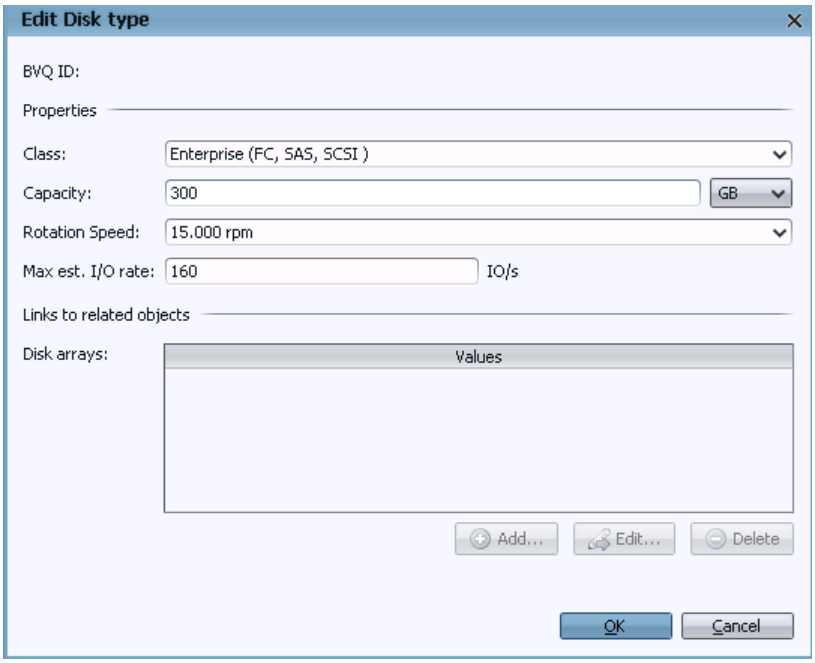

Qualifizierung Disk Controller

Für die Qualifizierung des Controllers wählen Sie die entsprechende Class und Room aus und tragen die Cache Size ein. Die Room Informationen muss eingetragen werden damit das Objekt (Controller) voll qualifiziert werden kann. Die Cache Size, MTM und Seriennummer sind aus dem Controller zu entnehmen.
Controller Cache Size ermitteln und als Enterprise, Midrange oder Lowrange klassifizieren.

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Controllers...* und wählen sie den zu qualifizierenden Controller aus.

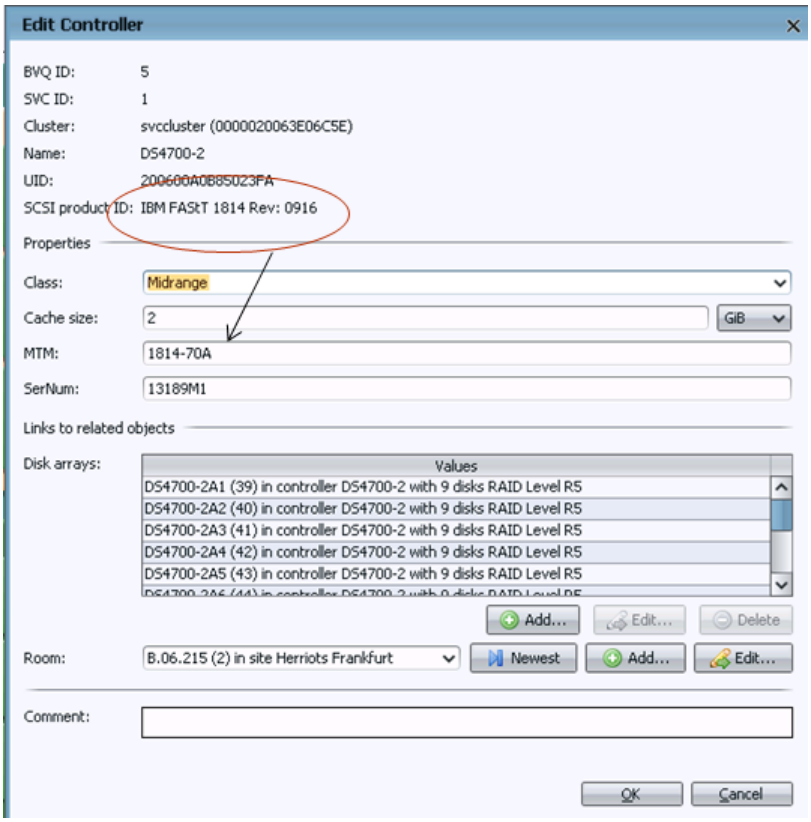

Qualifizierung Disk Arrays

In Disk Arrays werden die Performance Eigenschaften der darunter liegenden Hardware (MDisks) erfasst. Die erfassten Daten der Disk Arrays können später bei der Generierung von Reports genutzt werden (z.B. Tabelle aller Raid 5 Arrays).

Bei der Storwize V7000 werden Disk Arrays automatisch für jede einzelne MDisk angelegt. Es ist keine Benutzeraktion notwendig. Möchte man jedoch mehrere MDisks unter einem Disk Array zusammenfassten, so kann man die bestehenden Disk Array verwerfen und neue anlegen.

Tragen Sie Disk Array Name, Raid Level, Disktyp, Anzahl der Disks (ohne Hotspares) und den Controller ein.
Die Estimated I/O density (I/O Dichte) wird aus den Angaben des Raid Level und Disktyp (max est. I/O Rate) berechnet. Dieser Wert kann bei Abweichungen überschrieben werden.

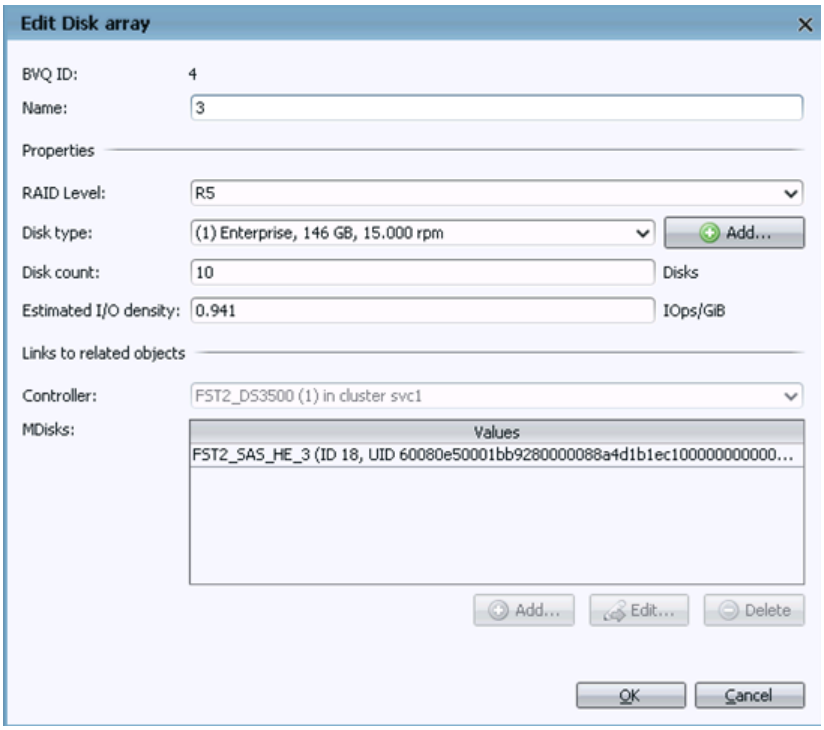

Nutzen Sie zur besseren Visualisierung von Controller, Disk Array und MDisk den unter Favoriten gespeicherten Treemap Pfad:
Favorites-> Predefined favorites -> Step by Step:Qualify and edit objects -> Define disk array for controllers.

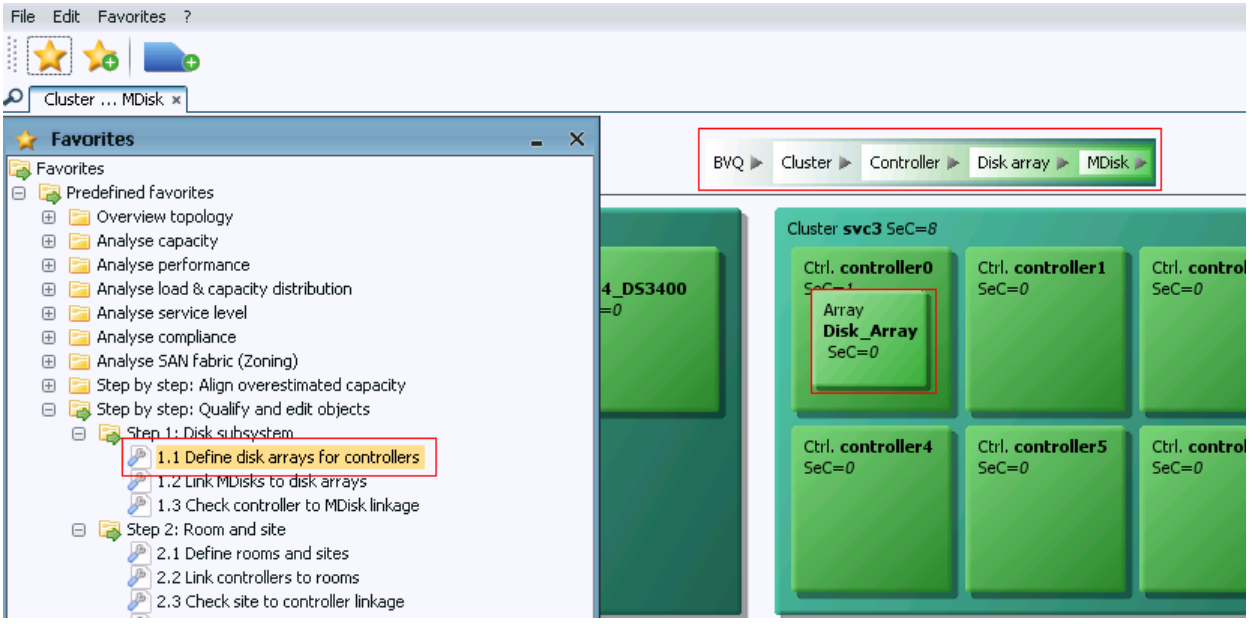

Wechseln Sie anschließend in den Treemap Pfad: BVQ -> Cluster -> Controller -> MDisk -> Disk array und wechseln Sie zur Object Color: BVQ qualified. MDisks die noch keinem Array zugeordnet sind werden daraufhin durch die Farbe Rot untermalt.

Sie können direkt in der Treemap, durch editieren des Controllers, ein neues Disk Array für den Controller erstellen und anschließen die MDisks dem neuem Array zuweisen.

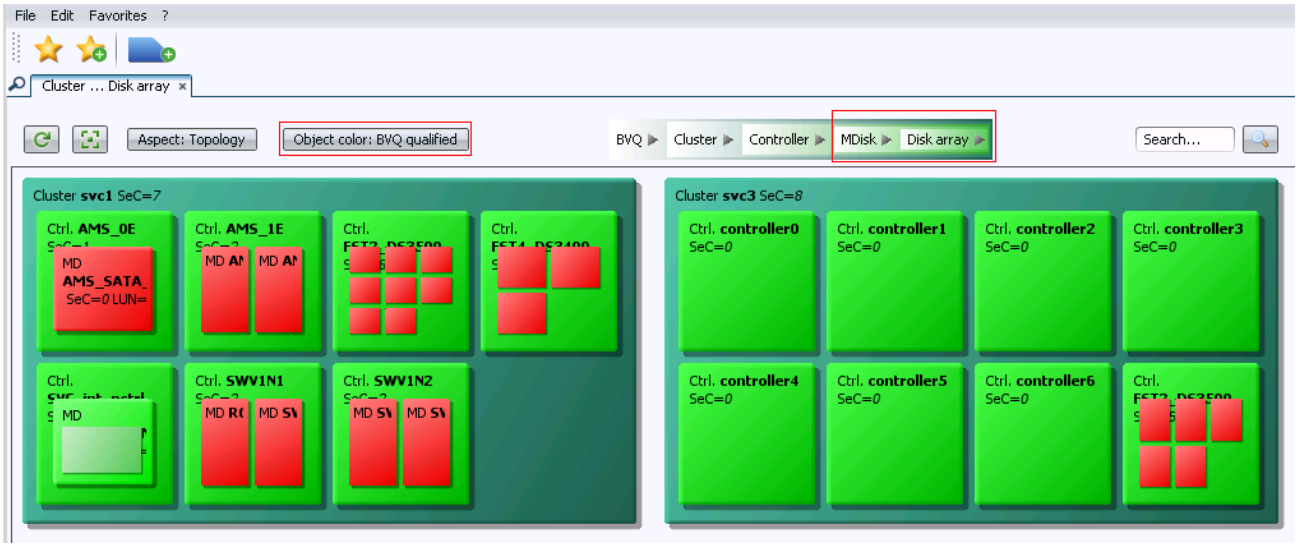

Qualifizierung Hosts

Die Hosts werden über den BVQ Scanvorgang erfasst. Für die Qualifizierung ist die Angabe des Rooms erforderlich. Zudem kann das Comment Feld (Kommentarfeld) für Kontakt o. A. verwendet werden.

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Hosts* und wählen sie den zu qualifizierenden Host aus.

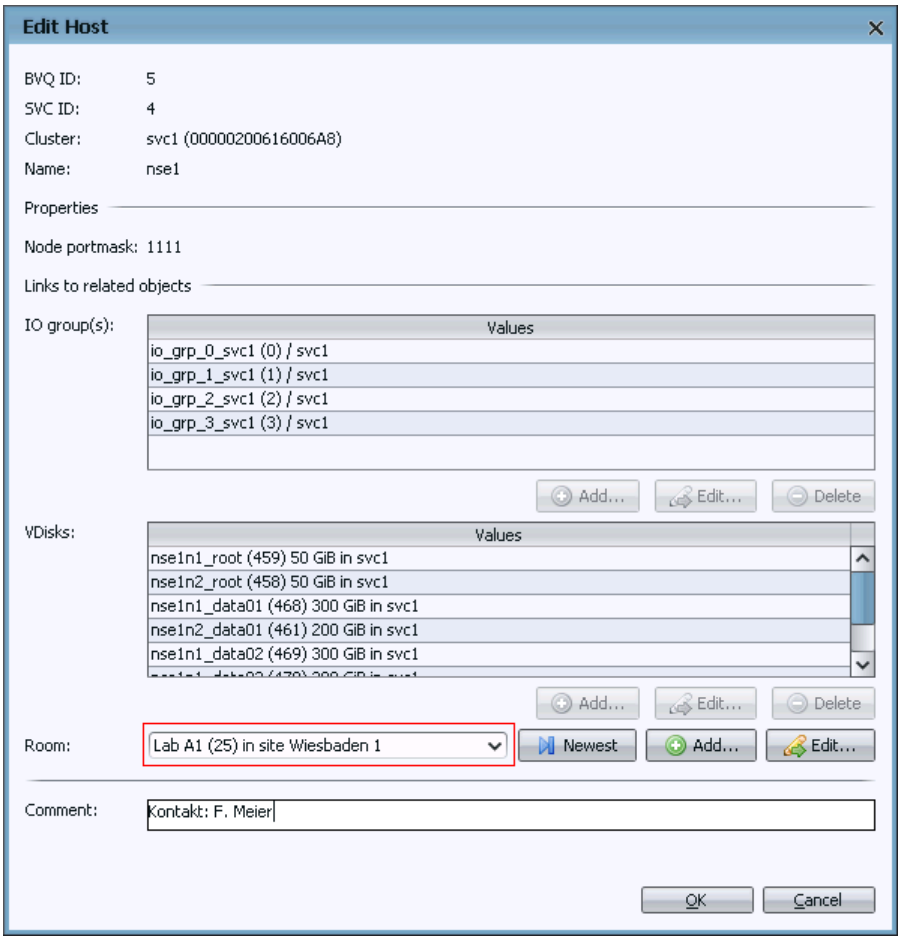

Qualifizierung SVC Nodes

Die Nodes werden auch über den BVQ Scanvorgang erfasst. Für die Qualifizierung ist die Angabe des Rooms erforderlich.

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Nodes* und wählen sie den zu qualifizierenden Node aus.

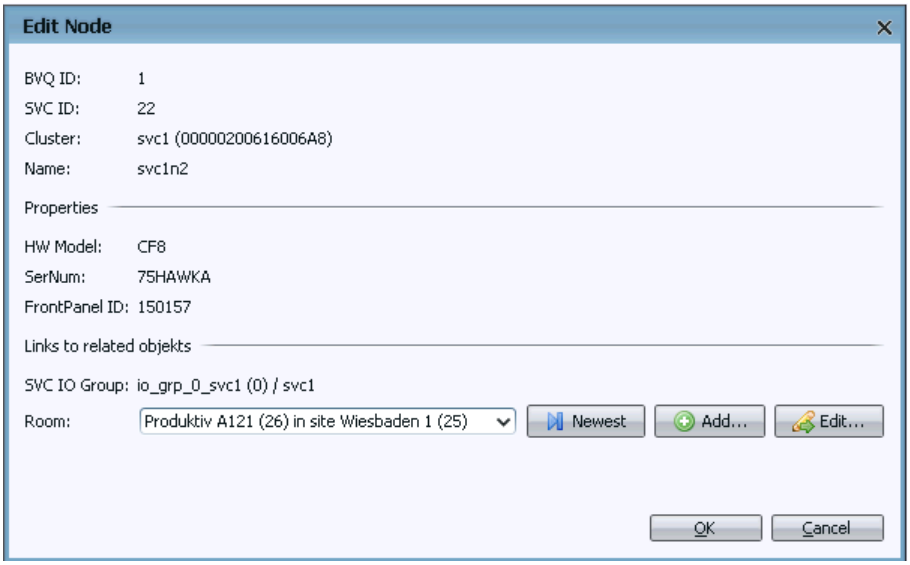

1.12.1.6 Qualifizierung BVQ Accounting Paket

Das Accounting Paket ist Bestandteil der Premium Edition und ist optional für die Performance Edition erhältlich. Führen Sie die in diesem Abschnitt beschriebenen Qualifizierungsschritte aus, wenn Sie im Besitz der Premium Edition oder der Performance Edition mit dem Accounting Paket sind.

Für das Accounting Paket werden Unternehmensinformationen benötigt, und dabei stellen sich die Fragen: "Welche Applikationen werden im Unternehmensumfeld eingesetzt ?" oder "Wie lassen sich diese Applikationen abrechnungstechnisch sinnvoll erfassen?"
Im BVQ ist es daher nötig die entsprechenden Informationen zusammenzutragen und anschließend sinnvoll zu verknüpfen.

- Was sind meine Kostenstellen?
  z.B. Kostenstelle: Personalabteilung, Kunde xy oder Software Entwicklung
  -> Diese Informationen werden in der Regel in die Gruppe "Costcenter" eingetragen.

- Welche Applikationsgruppen laufen unter den einzelnen Kostenstellen ?
  z.B. Gehaltsabrechnung, Zeiterfassung
  -> Diese Informationen werden in der Regel in die Gruppe "Application Group" eingetragen.

- Welche Applikationen laufen in den Applikationsgruppen ?
  z.B. SAP Prod, SAP Test, Notes, Citrix oder die Entwicklungssoftware
  -> Diese Informationen werden in der Regel in die Gruppe "Application" eingetragen.

- Sind die Applikationen auf verschieden Systemen verteilt ?
  z.B. Tablespaces, Backup
  -> Diese Informationen werden in der Regel in die Gruppe "VDisk Group" eingetragen.

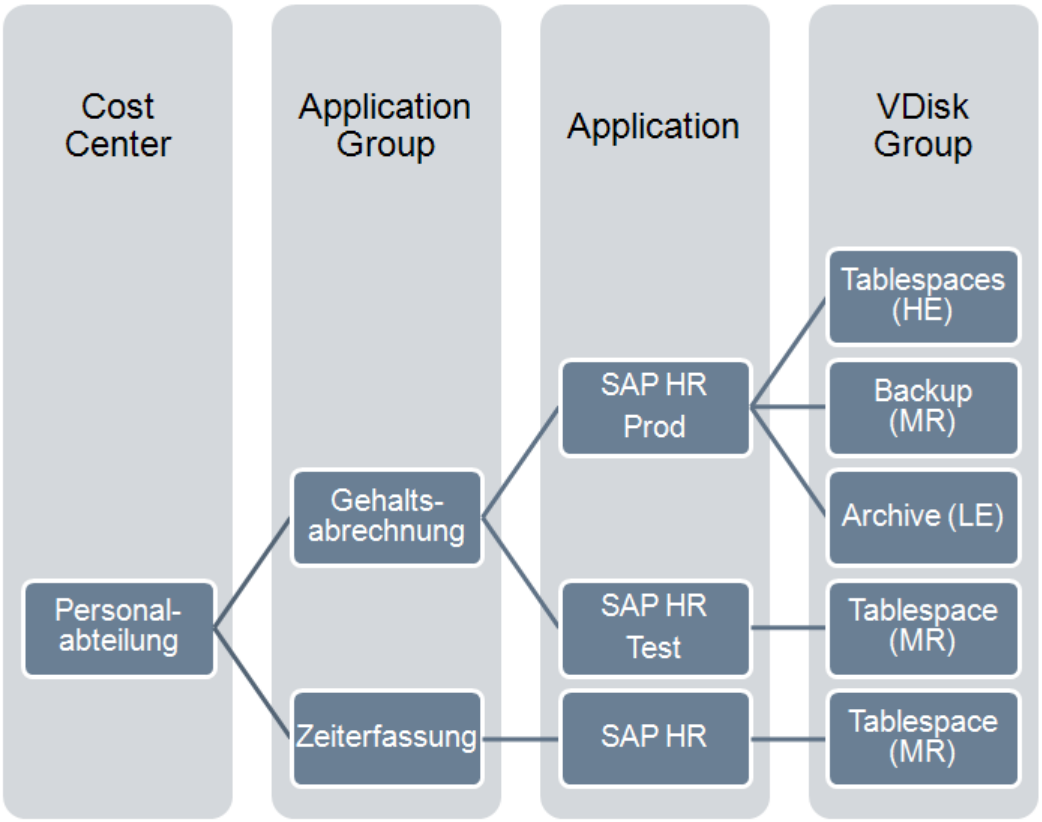

Die VDisk Gruppen stellen somit die Basis des Accounting Pakets dar. Mit den VDisk Gruppen werden später die SLA`s, Datenklassen, Mobility Zones, VDisks und die VDisk Gruppen Relationen verknüpft.

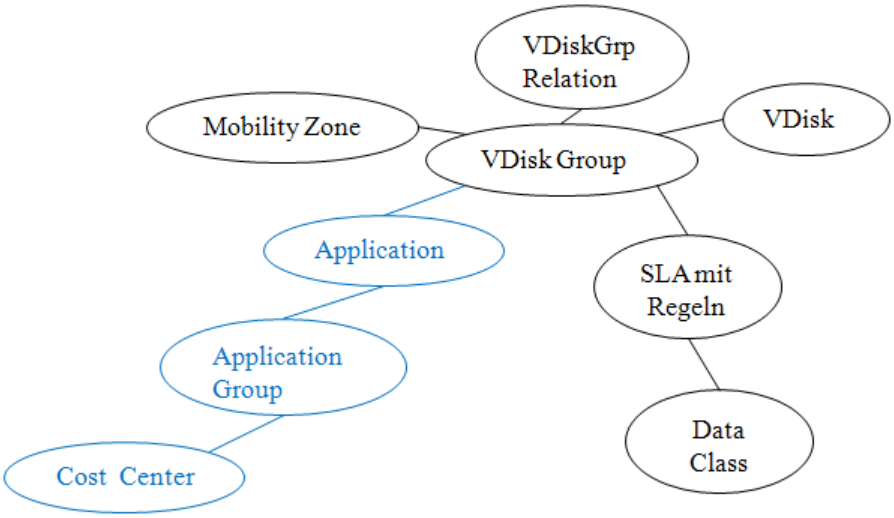

Qualifizierung Costcenter, ApplicationGroup, Application, VDiskGroup, VDisk

Die Qualifizierung gestaltet sich am sinnvollsten wenn man mit den Kostenstellen beginnt, und anschließend zu den Applikationgruppen, Applikationen und zuletzt die Qualifizierung der VDisk Gruppen durchführt. Dadurch kann das neue Objekt stets mit einem vorhandenen verbunden werden.

Objekt Cost Centers

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Cost centers...* und klicken sie auf Add um die verschiedenen Kostenstellen zu erstellen.

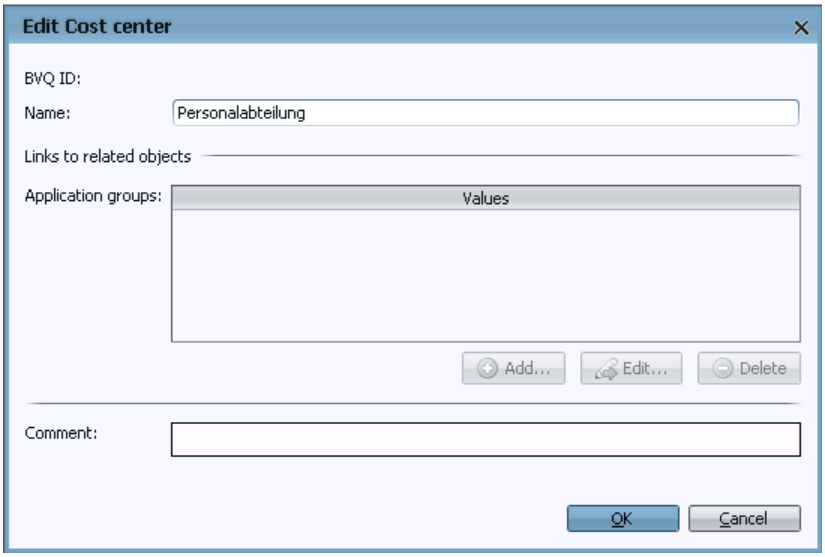

Objekt Application groups

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Application groups...* und klicken sie auf Add um die Applikationsgruppen zu erstellen. Weisen Sie jeder Applikationsgruppe die passende Kostenstelle zu.

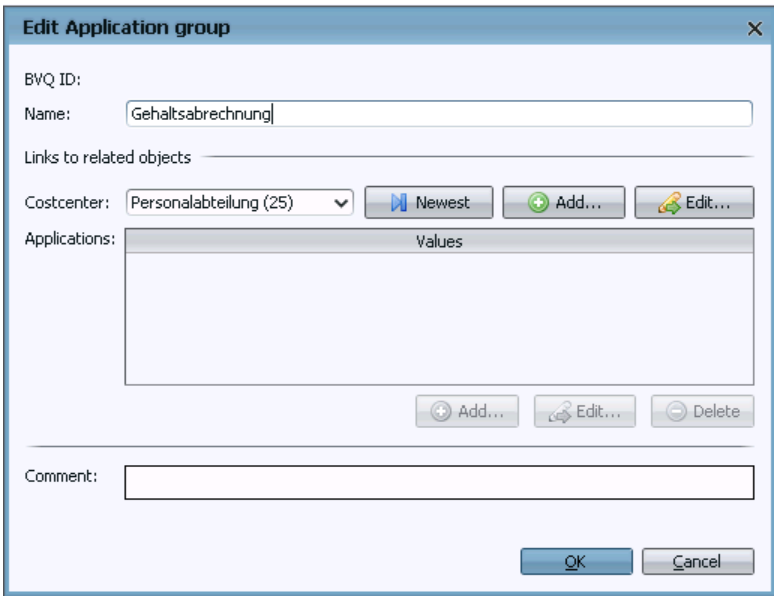

Objekt Application

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Applications...* und klicken sie auf Add um die Applikationen zu erstellen. Weisen Sie jeder Applikation die passende Applikationsgruppe zu.

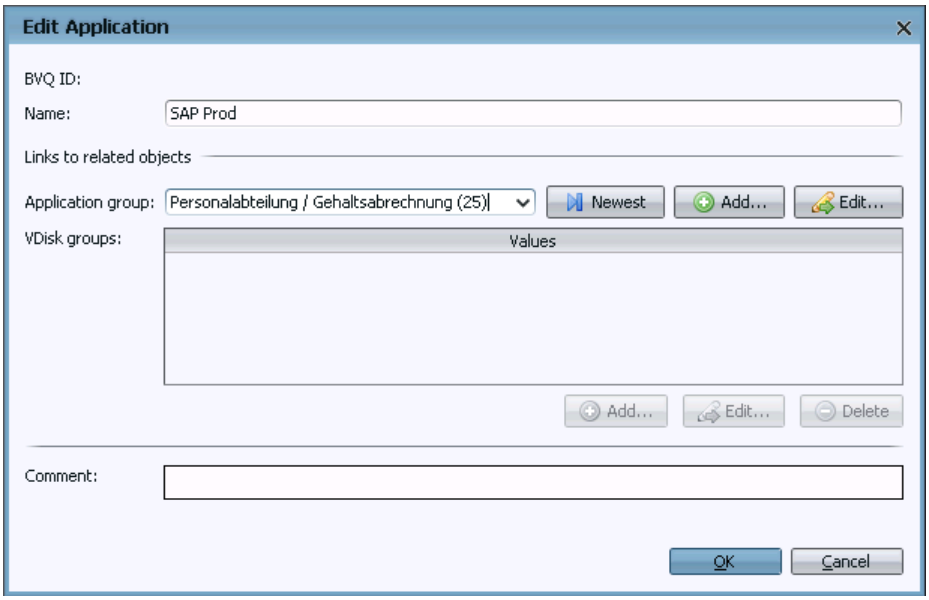

Objekt VDisk groups

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> VDisk groups...* und klicken sie auf Add um die VDisk Gruppen zu erstellen. Weisen Sie jeder VDisk Gruppe die passende Applikation zu.

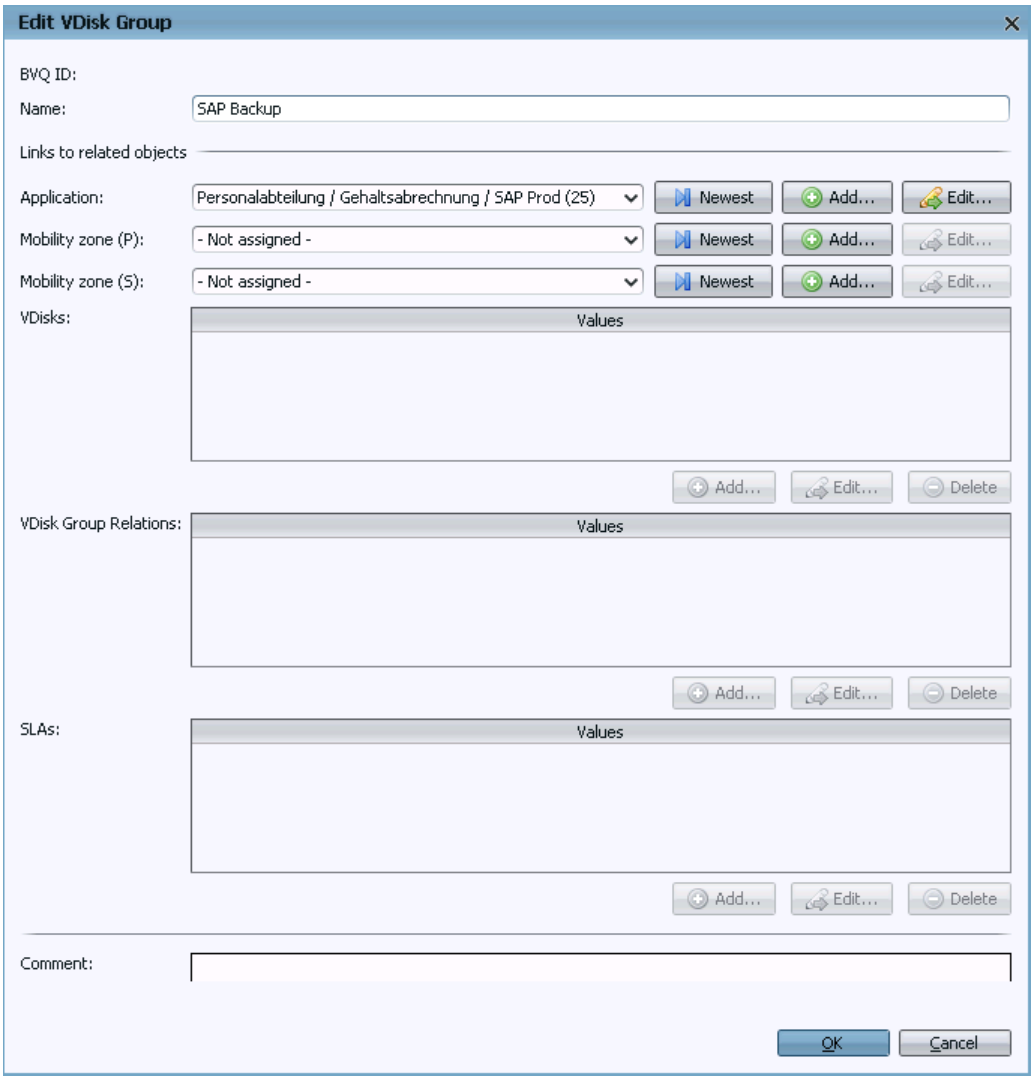

VDisk

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> VDisks...* und wählen eine VDisks aus, welche sich mit einer zuvor angelegten VDisk Gruppe passend verbinden lässt.
Nutzen Sie den Namensfilter um VDisks schnell zu finden.

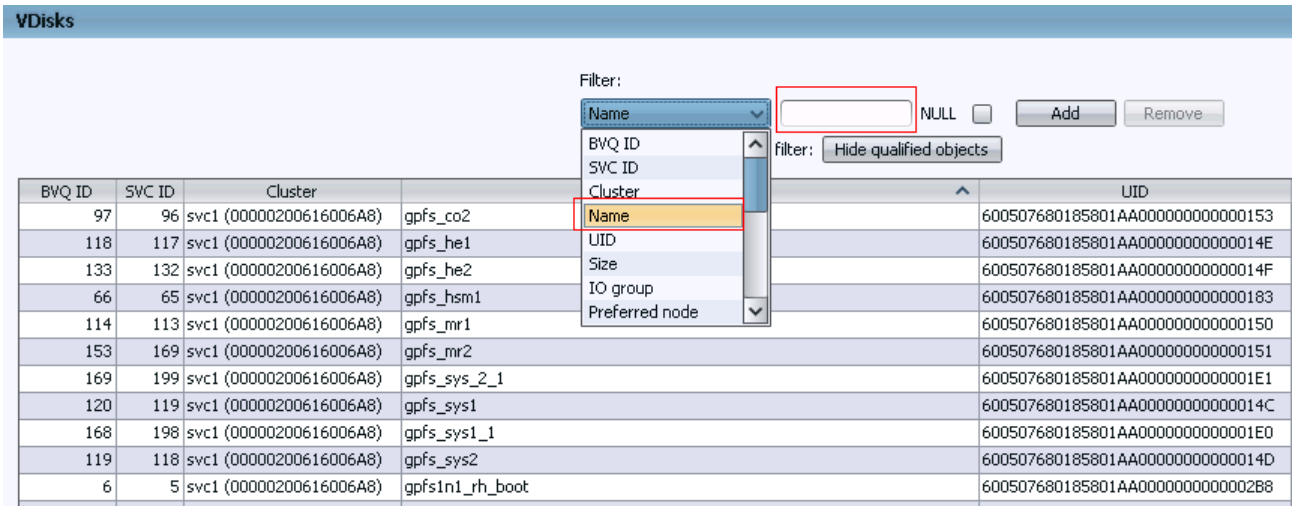

Weisen sie ihr daraufhin die gültige VDisk Gruppe zu. Führen Sie die Qualifizierung für die anderen VDisks ebenfalls durch.

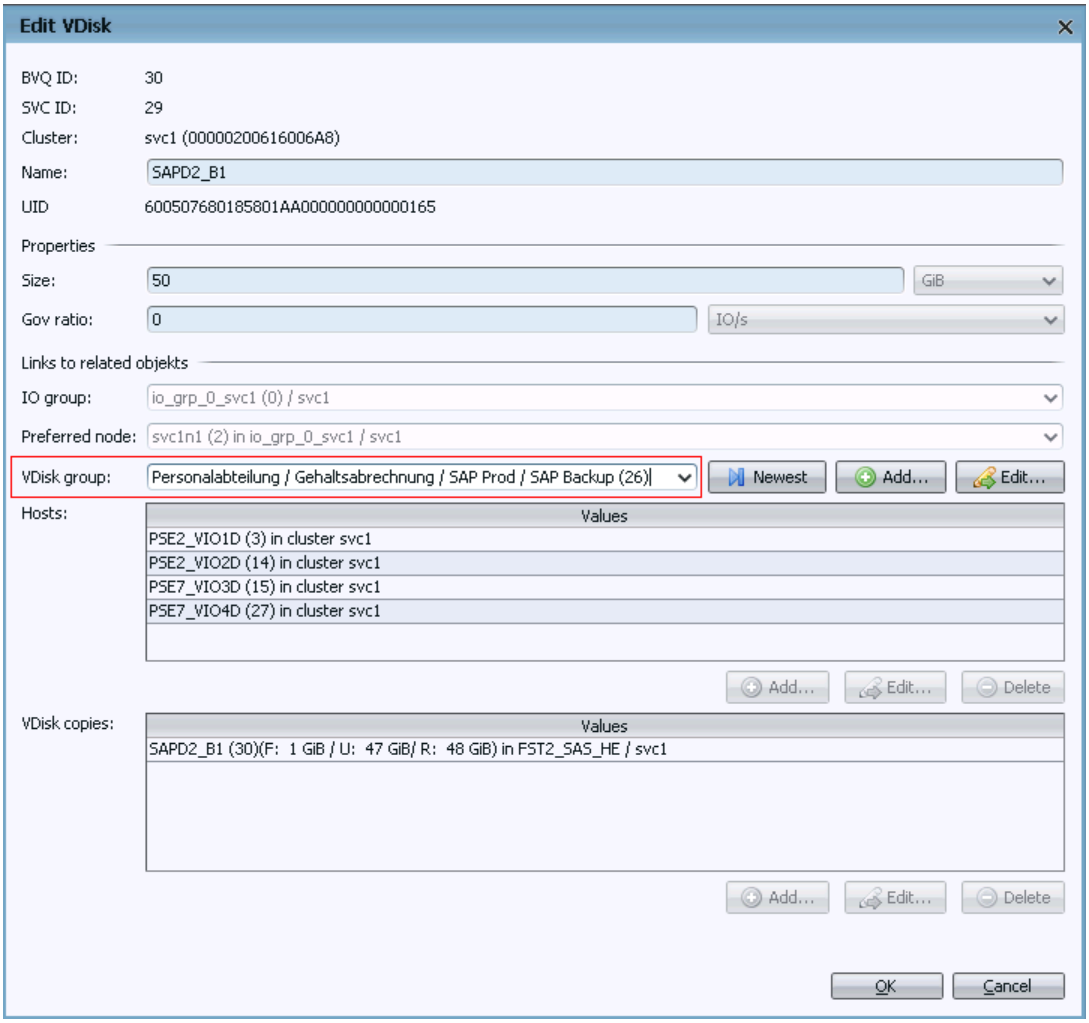

Nach Abschluss dieser Qualifizierungsschritte können die ersten Auswertungen im BVQ erfolgen (z.B. die Kapazitätssicht / Performancesicht auf eine Kostenstelle).

Kapazitätssicht:

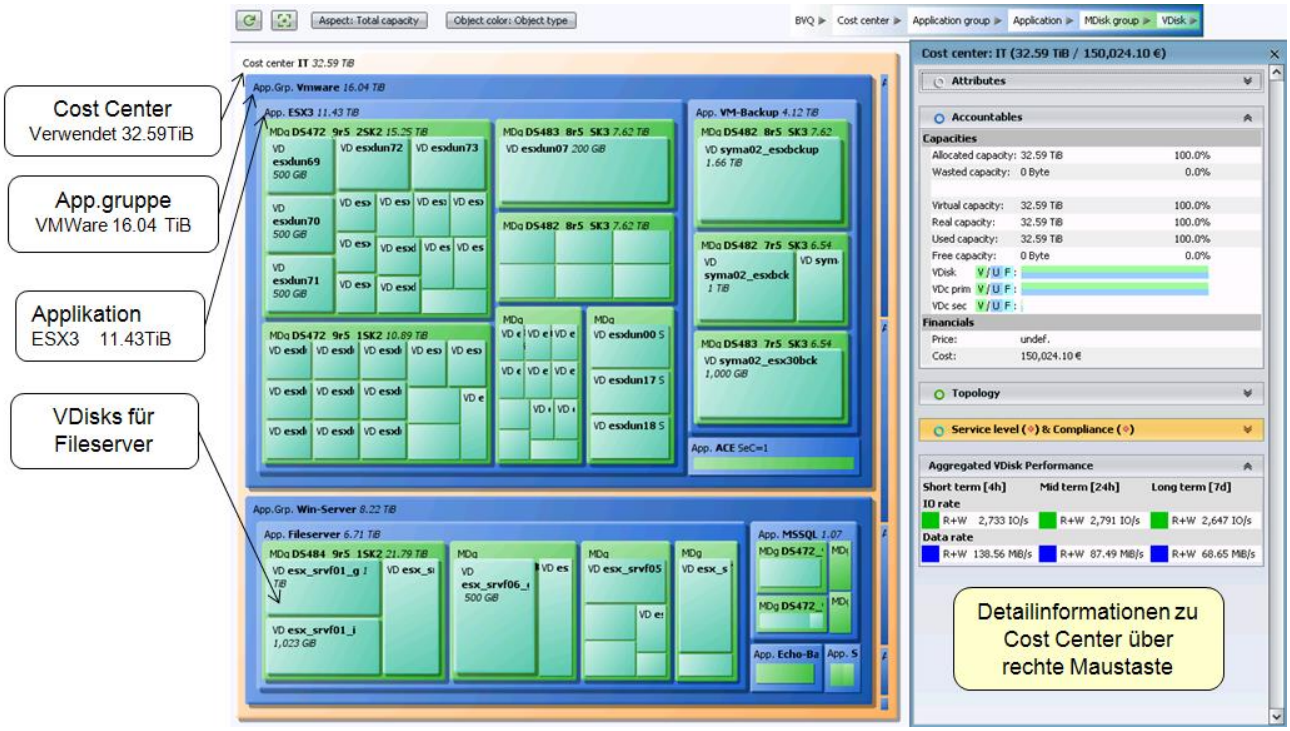

Performancesicht:

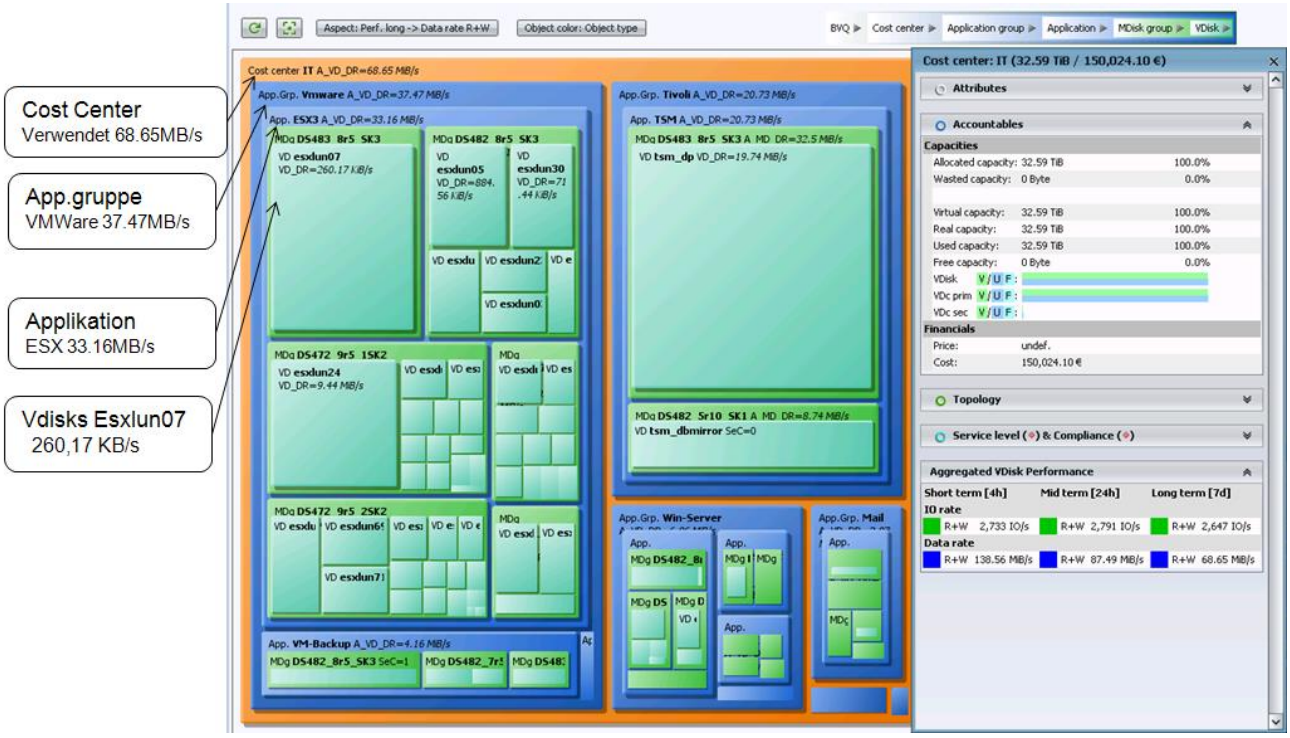

Qualifizierung MobilityZone

Beginnen Sie nun die Mobility Zones einzurichten. Mobility Zones legen fest, auf welchen Storage Pools die VDisks einer VDisk Gruppe platziert werden sollten. Somit können Regeln für Ablageorte von VDisks festgelegt werden, die dann beispielsweise auf Ihre Einhaltung hin überprüft werden können (VDisk X sollte nur auf den MDisk groups Y & Z abgelegt werden).

Erstellen und MDisk Gruppen zuweisen
Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Mobility zones...**und klicken Sie auf *Add**.

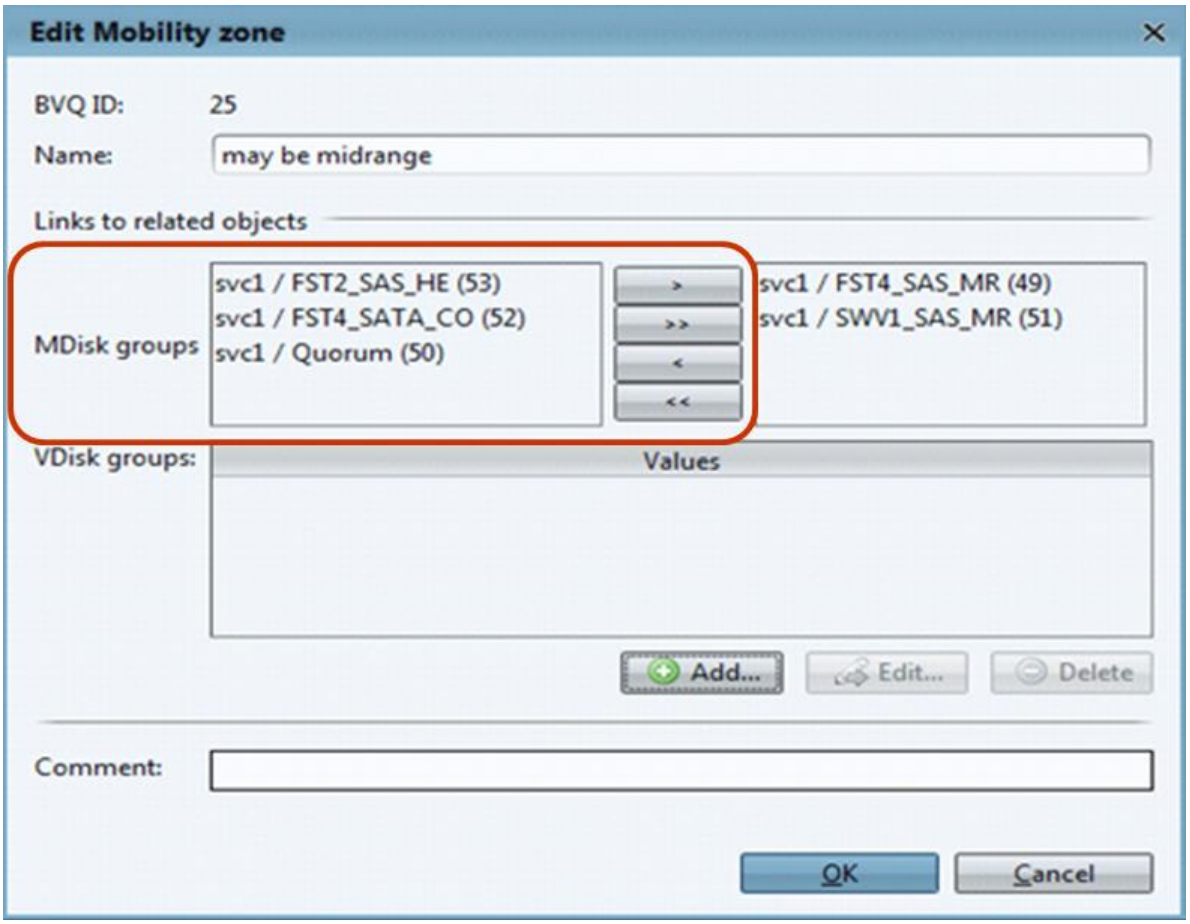

Mobility Zones den VDisk Gruppen zuweisen
Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> VDisk groups...* und wählen sie eine passende VDisk Gruppe aus.
Eine Mobility Zone kann hierbei für mehrere VDisk Gruppen verwendet werden.
Geben Sie für die Mobility Zone (P) die Mobilty Zone an, in welcher sich die primären Kopien der VDisks in der VDisk Gruppe befinden sollen.
Geben Sie für die Mobility Zone (S) die Mobilty Zone an, in welcher sich die sekundären Kopien der VDisks in der VDisk Gruppe befinden sollen. Somit können gespiegelte VDisks (VDisk Mirror) unterschiedlichen Mobility Zones zugewiesen werden.

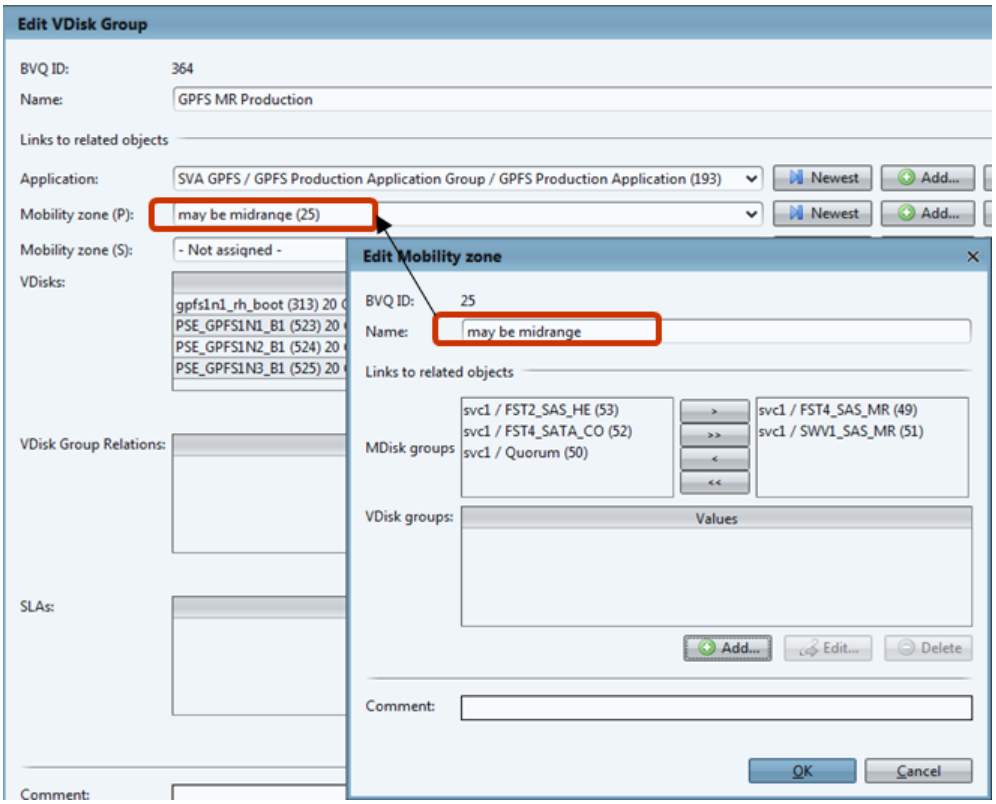

Mobility Zone Compliance Check

Über **Object Colors: Mobility zone compliance** werden in der Einfärbung die falsch platzierten Objekte rot angezeigt. Korrekt platzierte Objekte werden immer grün angezeigt. Blaue Objekte weisen darauf hin, dass keine Complianceregeln bestehen.

Nutzen sie auch hier die Step by Step -> Compliance rules Funktion um schnell einen Überblick über die Compliance zu erhalten.

Qualifizierung VDiskGroupRelation

VDisk Group Relations ermöglichen es sicherzustellen, dass alle VDisks aus einer VDisk Group speziellen Platzierungsregeln genügen z.B. zwei über LVM gespiegelte Bereiche dürfen nicht am gleichen Ort gespeichert werden.
Die VDisk Gruppen werden dazu untereinander auf ihre Platzierung geprüft.
Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> VDisk group relations...*und klicken Sie auf Add.
Wählen Sie die Placement behavior und den Objekt Typ aus. Die Objekttypen umfassen dabei MDiskgroups, MDisks, Arrays, Controllers, Nodes, IOGroups, Cluster, Rooms und Sites.

- **Antilocated:** dürfen nicht auf dem gleichen Objekt wie die jeweils andere Disk Gruppe liegen.
- **Collocated:** müssen auf dem gleichen Objekt wie die jeweils andere Disk Gruppe liegen.

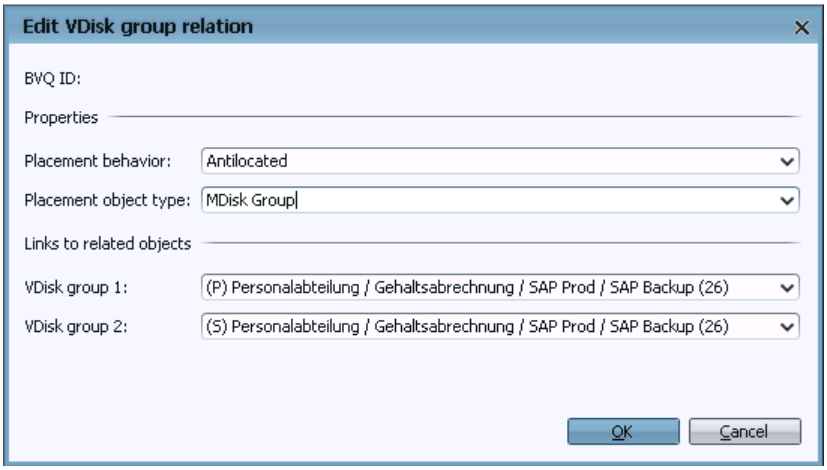

1.12.1.7 Qualifizierung BVQ SLM Paket

Das Service-Level-Management Paket ist Bestandteil der Premium Edition und optional für die Performance Edition erhältlich.

BVQ SLA

- **Ein SLA in BVQ enthält Preis/GiB, Zeitraum und Data Placement Regeln.**
  Beispiel:
  Vom 01.01.2011 bis 31.12.2011 werden die SLA beschriebenen VDisks und ihre Kopien auf der Datenklasse A gespeichert. Für diesen Service werden xx€/GiB berechnet.

BVQ Datenklasse & BVQ Speicherklasse

BVQ fördert die Platzierung von klassifizierten Datenbereichen in einer passenden und möglichst günstigen Speicherklasse.

Datenklassen (Data classes): Kundenorientiert mit Speicherpreis €/GiB
Speicherklassen (Storage classes): Anbieterorientiert mit Speicherkosten €/GiB

BVQ Speicherklasse

Die BVQ Speicherklasse stellt die Klassifizierung vorhandener Speicher-Kapazitäten dar.

Mit ihr sind die Kosten (BVQ „Kosten" in €/GiB) verbunden die beim Betreiber anfallen.

Die folgende Tabelle stellt eine mögliche Speicherklassenaufteilung dar. Anzahl und die Benennung der Speicherklassen sind nicht fest vorgeschrieben.

Speicherklasse	Beschreibung	Kosten der Speicherklasse (fiktiv)
Gold	schnellste Speicherklasse	75
Silber	mittlere Speicherklasse	40
Bronze	langsamste Speicherklasse	15

BVQ Datenklasse

Die BVQ Datenklasse stellt die geforderte Klassifizierung von Datenbereichen (Volumes) dar.

An sie ist der Preis (BVQ "Abrechnungspreis" in €/GiB) geknüpft, den der Anbieter vom Endkunden verlangt.

Die folgende Tabelle stellt ein mögliches Datenklassenszenario dar. Anzahl und Benennung der Datenklassen sind nicht fest vorgeschrieben.

Daten-klasse	Beschreibung	Speicherklasse Bronze (fiktive Kosten = 15)	Speicherklasse Silber (fiktive Kosten = 40)	Speicherklasse Gold (fiktive Kosten = 75)	"Abrechnungspreis" (fiktiv)
A+	Datenklasse mit den höchsten Anforderungen	★ (inakzeptabel)	★ (inakzeptabel)	★ (optimal)	100
A		★ (inakzeptabel)	★ (akzeptabel)	★ (optimal)	90
B+		★ (inakzeptabel)	★ (optimal)	★ (akzeptabel)	75
B		★ (inakzeptabel)	★ (optimal)	★ (inakzeptabel)	50
C+		★ (optimal)	★ (akzeptabel)	★ (inakzeptabel)	40
C	Datenklasse mit niedrigsten Anforderungen	★ (optimal)	★ (inakzeptabel)	★ (inakzeptabel)	20

BVQ Platzierungsbewertung

In BVQ kann die Bewertung der Platzierung für alle Kombinationen aus Daten- und Speicherklassen Kundenspezifisch eingegeben werden.

Die Platzierung von Datenbereichen auf einer gleichwertigen Speicherklasse ist i.d.R. bevorzugt.

Die Platzierung von hochwertigen Datenbereichen auf niederwertigen Speicherklassen ist i.d.R. riskant aber günstig.

Die Platzierung von niederwertigen Datenbereichen auf hochwertigen Speicherklassen ist i.d.R. zu teuer.

Qualifizierung Storage classes

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Storage classes...* und klicken Sie auf *Add*.

- Wählen Sie einen passenden Namen (z.B. Gold, Silber, Bronze)
- Wählen Sie für die visuelle Darstellung eine Farbe aus.
- Tragen Sie den €/GiB Preis ein.
- Die MDisk Gruppen müssen im Anschluss über *Edit -> MDisk groups.. hinzugefügt werden. Aktuell ist es in dem Auswahldialog Edit->Storage class noch nicht möglich die MDisk Gruppen direkt auszuwählen.
- Das "Data class to Storage class placement rating" wird später hinzugefügt, da diese erst angelegt werden müssen.
- Wiederholen Sie die Schritte für die anderen Speicherklassen.

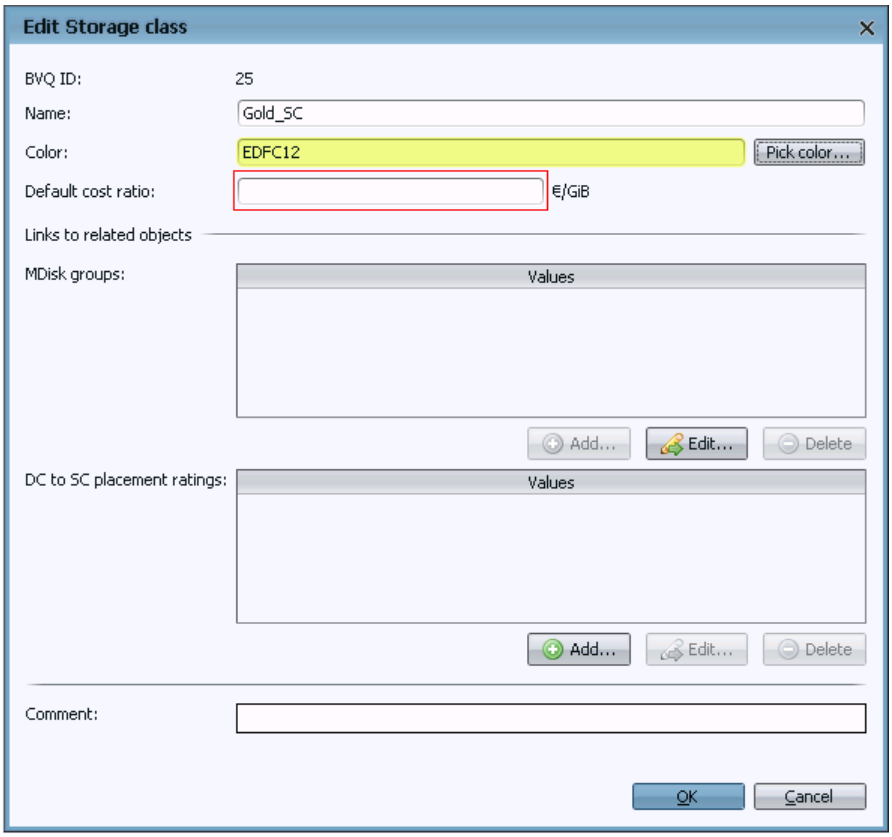

Qualifizierung Data classes

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> Data classes...* und klicken Sie auf *Add*.

- Wählen Sie einen passenden Namen (z.B. A+, A, B+ ..)
- Wählen Sie für die visuelle Darstellung eine Farbe aus.

- Tragen Sie den €/GiB Preis ein.
- Das "Data class to Storage class placement rating" und "Active SLA`s" werden später hinzugefügt, da diese erst angelegt werden müssen.
- Wiederholen Sie die Schritte für die anderen Datenklassen.

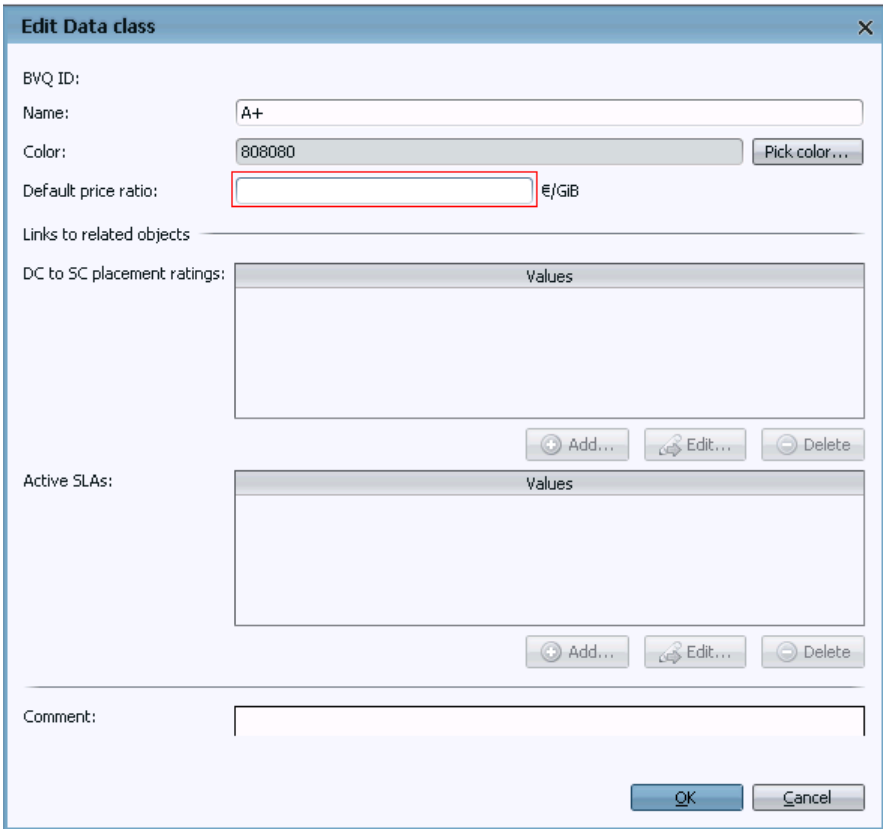

Qualifizierung DC to SC placement ratings

Erstellen Sie nun eine gültige Matrix aus Daten und Speicherklassen (Siehe Beispiel: #Datenklasse).

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> DC to SC placement ratings...* und klicken Sie auf *Add*.

- Wählen Sie eine Datenklasse aus.
- Wählen Sie eine Speicherklasse aus.
- Verknüpfen Sie diese anschließend über das **Placement rating:** mit einer gültigen Klassifizierung (Optimal, Suboptimal, Unacceptable).

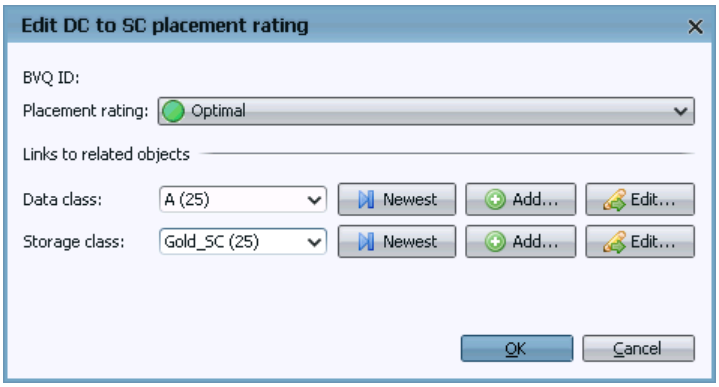

Achten sie darauf, dass für alle möglichen Kombinationen ein DC to SC rating vorliegt. Dadurch können im BVQ falsch platzierte Daten schnell erfasst und korrigiert werden.

Qualifizierung SLA

- SLA`s werden immer an VDisk Gruppen gebunden.
- Sie legen den Zeitraum fest, auf welcher Datenklasse gespeichert wird und was dieser Service kostet.
- Die Einhaltung der SLA`s kann über *Object colors:Overall Compliance geprüft werden.
- Secondary copy data class beschreibt auf welcher auf welchem Speicher die VDisk copies abzulegen sind.

Starten Sie die Qualifizierung in der Menüleiste mit *Edit -> SLA`s...* und klicken Sie auf *Add*.

- Geben Sie einen Zeitraum für die Gültigkeit des SLA`s an.
- Tragen Sie den €/GiB Preis ein.
- Betroffene VDisk Gruppe auswählen.
- Wählen Sie die Datenklasse für die Primary und Secondary Copy aus.
  Eine VDisk ist erst vollqualifiziert wenn sie mit einem SLA verknüpft ist.

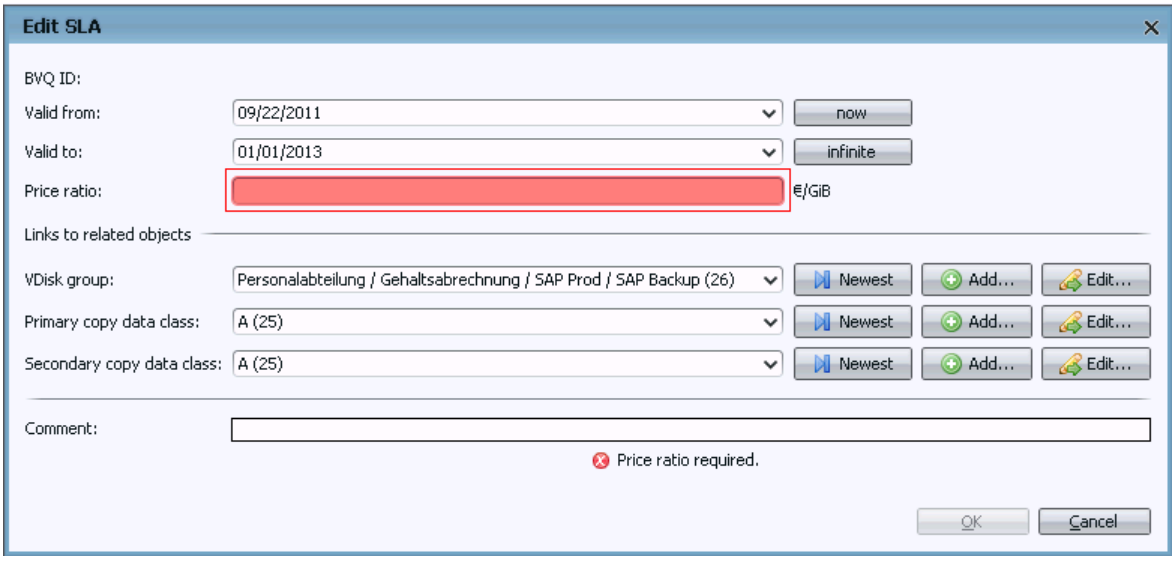

1.12.1.8 Qualifizierungs-Abhängigkeiten - BVQ Qualified

Die folgende Tabelle gibt Aufschluss über die Qualifizierungsabhängigkeiten der einzelnen BVQ Objekte. Ist die Qualifizierungsabhängigkeit erfüllt, dann wird das Objekt als BVQ qualifiziert angezeigt, und die damit verbundenen Informationen können für spätere Auswertungen im BVQ bzw. im Reporting genutzt werden.

BVQ-Objekt	Entry Edition	Performance Edition benötigt:	Perf. + Acc. Edition benötigt:	Perf. + SLM Edition benötigt:	Premium Edition benötigt:
VDiskGroups	⊖	⊖	✓	SLA	SLA
VDisks	✓	✓	VDiskGroup	VDiskGroup	VDiskGroup
VDiskCopies	ⓘ	ⓘ	ⓘ	ⓘ	ⓘ
MDiskGroups	✓	✓	✓	ResponseTime & StorageClass	ResponseTime & StorageClass
MDisks	Array	Array	Array	Array	Array
Nodes	✓	Room	Room	Room	Room
Hosts	✓	Room	Room	Room	Room
Controllers	✓	Room	Room	Room	Room
DiskArrays	ⓘ	ⓘ	ⓘ	ⓘ	ⓘ

BVQ-Objekt	Entry Edition	Performance Edition benötigt:	Perf. + Acc. Edition benötigt:	Perf. + SLM Edition benötigt:	Premium Edition benötigt:
DiskTypes	ℹ	ℹ	ℹ	ℹ	ℹ
Applications	–	–	ApplicationGroup	–	ApplicationGroup
Application Groups	–	–	Costcenter	–	Costcenter
Costcenter	–	–	ℹ	–	ℹ
SLA`s	–	–	–	PriceRatio,VDiskGroup, Primary- & Secondary copy data class	PriceRatio,VDiskGroup, Primary- & Secondary copy data class
Data classes	–	–	–	Eine Data class muss zu allen bestehenden Storage classes ein DC to SC placement rating besitzen	Eine Data class muss zu allen bestehenden Storage classes ein DC to SC placement rating besitzen
Storage classes	–	–	–	Eine Storage class muss zu allen bestehenden Data classes ein DC to SC placement rating besitzen	Eine Storage class muss zu allen bestehenden Data classes ein DC to SC placement rating besitzen
DC to SC placement ratings	–	–	–	PlacementRating + Dataclass + Storageclass	
Mobility zones	–	–	–	MDiskGroup	MDiskGroup
VDiskgroup relations	–	–	–	2 VDiskGroups	2 VDiskGroups
Rooms	–	Site	Site	Site	Site
Site	–	ℹ	ℹ	ℹ	ℹ

Legende:

	Nicht Bestandteil der Edition
	Keine Qualifizierung möglich /notwendig
	Es wird kein Objekt benötigt (BVQ Qualified Status liegt bereits vor)

1.12.1.9 Anzeige von Qualification, Compliance und Placement Rating

Mit den Editoren für VDisk group relations, Mobility Zones und SLAs ist die Abbildung von Compliance-Vorgaben, SLA-Regeln und Objekt Qualifizierung möglich. So lassen sich Zustände an verschiedenen Stellen der BVQ GUI anzeigen und schaffen dadurch einen erheblichen Mehrwert und Übersicht.

Folgende Möglichkeiten ergeben sich daraus:

Compliance Check – Die Einhaltung der eingegebenen Vorgaben prüfen.

Dies betrifft:

- **Data class compliance** Aktuelle VDisk Platzierung bezüglich ihrer Data Class
- **Mobility zone compliance** Aktuelle VDisk Platzierung bezüglich ihrer Mobility zones
- **VDisk group relation compliance** Aktuelle VDisk Platzierung bezüglich ihrer VDisk group relations
- **Overall compliance** Zusammenfassung der drei obigen Zustände

Placement rating – Eine Bewertung aus Data Class und Storage Class prüfen.

- Aktuelle Bewertung der VDisk Platzierung bezüglich ihrer Data Class und Storage Class

Qualification state – Die Vollständigkeit der BVQ Qualifizierung von Objekten prüfen.

In der GUI gibt es grundsätzlich zwei verschiedene Darstellung die sich zum Anzeigen der Zustände eignen:

- Treemap
- List views

Mögliche Zustände der Compliance

Tabelle 77 Compliance Zustände

Symbol	Zustand	Erklärung
	Compliant	Aktuelle Konfiguration des Objekts erfüllt alle vorhandenen Vorgabe(n)
	Not Compliant	Aktuelle Konfiguration des Objekts erfüllt nicht alle vorhandenen Vorgabe(n)
	Undefined	Für das Objekt sind keine Vorgaben definiert
	Not	Objekt (-Typ) ist nicht betroffen

	Applicable	

1.12.1.10 Activation and usage of BVQ CLI

The BVQ CLI is hidden and has to be activated. To activate the CLI insert "BVQ_CLI" into the search field and press the search button (magnifying glass).

The CLI is now ready to use (orientation via help menu)

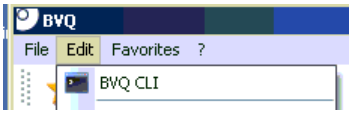

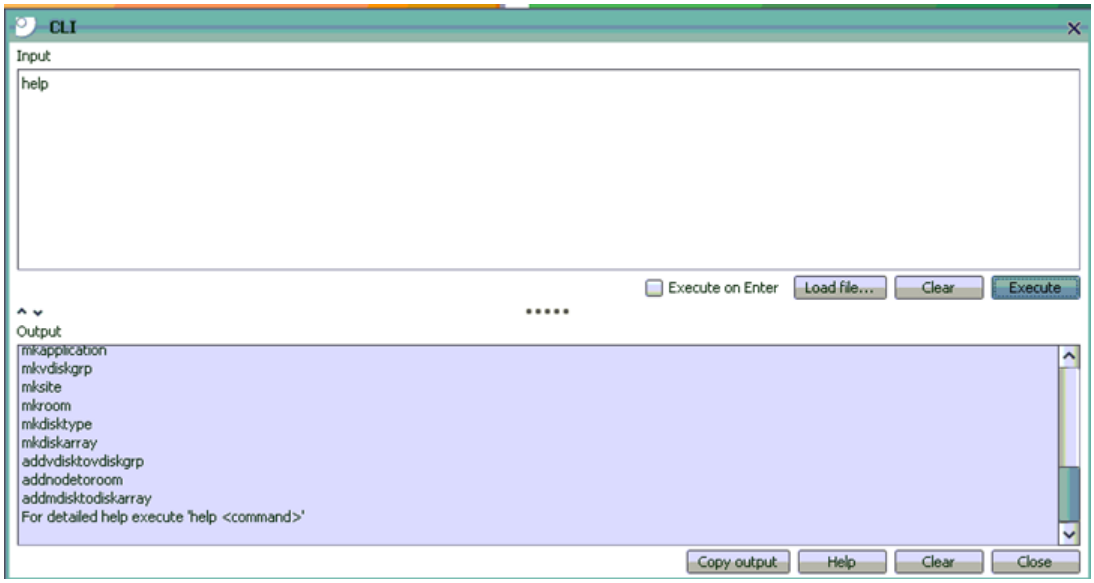

1.12.1.11 Combining multiple SVC controller to one BVQ controller (V7 und HDS)

In particular for HDS and Storwize V7000.
To be performed before arrays are created for the controller objects.

1. Create new BVQ controller with a suitable name

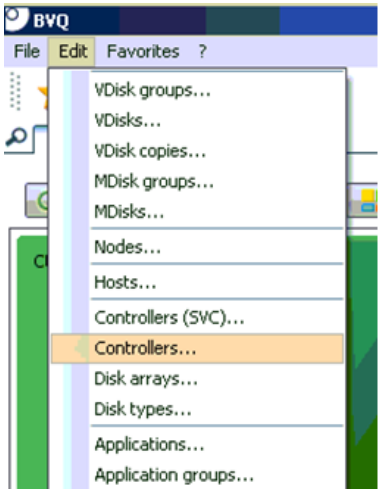

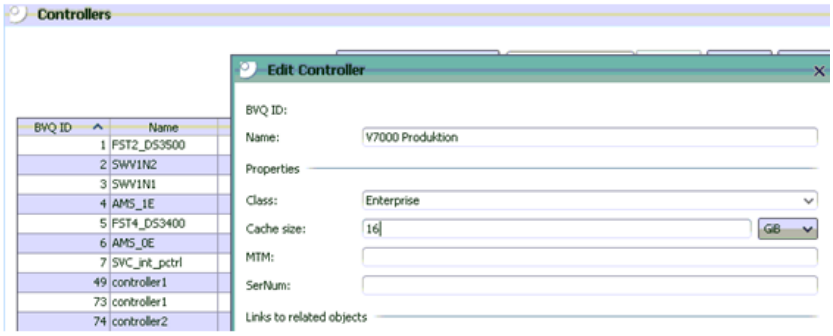

2. Connect the SVC controllers to a new BVQ controller - currently only possible via the treemap

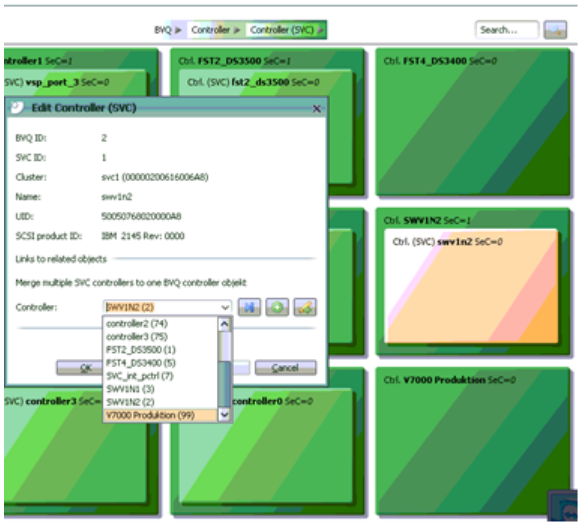

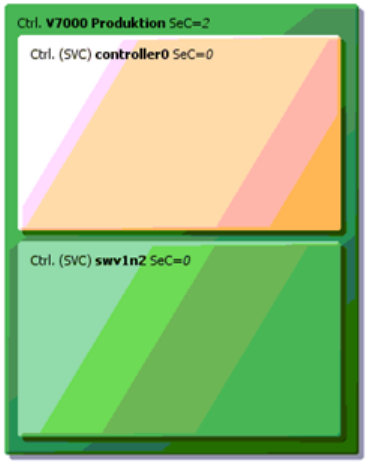

1.12.1.12 Faster creation of disk arrays by using BVQ CLI

The definition of disk arrays can get very complex, if done using the GUI. It is much faster done by using the BVQ CLI.

1. Activate CLI: activation and usage of BVQ CLI
2. Preparation: create an array with GUI to determine the correct IO density
3. Read controllerID und disk type ID
4. Specify the raid factor with R5 and not 5 in the command
5. Put everything together in Notepad and paste it into the command field (maybe keeping it for documentation)

```
> help mkdiskarray
usage: mkdiskarray [options]
 -controller <contoller id>   id of the controller to associate this disk array
with
 -diskcount <number of disks>   number of disks used in this disk array
 -disktype <disk type id>   id of the disk type to associate this disk array
with
 -iodensity <estimated IO density in IOps/GiB>
 -name <disk array name>   name of the disk array
 -raidlevel <raidlevel>   [R0,R1,R4,R5,R6,R10] RAID level of the disk array

Combined editing and then copy and paste
mkdiskarray –controller 15 –diskcount 10 –disktype 1 –iodensity 0.45 –name
Array01 –raidlevel R5
mkdiskarray –controller 15 –diskcount 10 –disktype 1 –iodensity 0.45 –name
Array02 –raidlevel R5
mkdiskarray –controller 15 –diskcount 10 –disktype 1 –iodensity 0.45 –name
Array03 –raidlevel R5
mkdiskarray –controller 15 –diskcount 10 –disktype 1 –iodensity 0.45 –name
Array04 –raidlevel R5
```

1.12.1.13 Qualification

Qualification

Jumping to the Ending of the der Metadata

- Anhänge:28
- Hinzugefügt von Norman Eichelbaum, zuletzt bearbeitet von Stephanus Abraham Le Roux am 24 Jul, 2013 (Änderung anzeigen)

To the Beginning of the MetadataIn this chapter the qualification of the Performance Edition and the single Extension-Packages is described.

In the Premium Edition, several qualification-steps, that are described here, have to be carried out, because both Extension-Packages are part of these steps.

In the Entry Edition, the qualification of the "MDisk to Array" is solely possible. See for that: http://wiki.sva.de/#Qualifizierung-QualifizierungDiskArrays

For that matter, no further attention is given to the Entry Edition in this chapter.

Indication of unqualified Objects

One obtains an overview about the current qualification-status through the Treemaps of the GUI.

1. Placing Object Colors on BVQ qualified.
2. All Treemap Objects, that are not yet qualified, are now colored with red.

3. Select the desired object i.e. an object-combination in the Treemap Navigator.
4. The unqualified object can be editted/qualified subsequently with a right-click.

Use the **Step-by-Step** that is stored in **Favorites: Qualify and edit objects Procedure for inspecting the Qualification-Progress** or to directly carry out the qualification in single Treemaps. The sequence of the qualification might (if necessary) slightly deviate from what is told in the User Manual Guide. It is important that valid combinations of the BVQ exist.

Qualifying Performance Edition

In the first step of the qualification, the topology-objects are qualified. These comply with the complete qualification-range in the Performance Edition and are part of the qualification-range of the Premium Edition.

For the first step at the qualification, the following information is needed.

- **Sites with Rooms**
  From all Hosts and Storage systems that have to be qualified, the Site and Room Information must be available.
  Where are the Storage Systems, SVC Knots and Hosts located?
  (Location and Room)
- **Storage System Information**
  Evaluation (Lowrange, Midrange, Highend), Cachesize, MTM Information (specific technical description) and Serial Number.
- **Disk Type**
  Which disk types are applied (SATA, SAS,SSD, 10kRPM, 15kRPM, max. IO\s)?
- **Disk Array Information**
  Reading out all existing Disk Arrays of the Storage Systems: Raid Level, Disk Type, amount of disks, Controller, MDisks
- **Determining Controller** Cache Size and classifying it as Enterprise, Midrange or Lowrange.
- **Server / Hosts:** Location and Contact
- **SVC Nodes:** Location

Hint: Discovery of technical Information

- Example: DS3500
  - Opening DS Storage Manager
  - Calling up Profile

- Reading out Information from Sections (Controller Cache, MTM, Serial Numbers, Raid Information)

Qualifying Site & Room

1. Start the Qualification in the Task Menu with *Edit -> Site*and klicken on *Add*.

2. Enter a Name for the Site. Repeat Procedure with every Site.

3. Carry out the same steps for the Sie Room (**Edit->Room**) and allocate to each Room the appropriate Site.

Clue: The Comment-Field can be utilized for additional Information, like e.g. Personal Contacts.

Qualifying Disk Types

For the Qualification of Disk Types, the disk-attributes have to be determined first.

- **Class:** Does it concern Lowrange (ATA, SATA, PATA), Enterprise (FC, SAS, SCSI) or SSD Storage ?
- **Capacity:** What capacity do the Disks have available?
- **Rotation Speed:** What numbers of revolution do the Disks have?
- **Max est. I/O rate:** Is the Speed determined out of rotation? This Value can also be adapted.

Starten the Qualificaton in the Task Menu with *Edit -> Disk types*and click on *Add*.

Qualifying Disk Controller

For the Qualification of the Controller you select the correspondent Class and Room and register the Cache Size. The Room Information must be entered so that the object (Controller) can be fully qualified. The Cache Size, MTM and Serial Number are to be gathered from the Controller. Determining Controller Cache Size and classifying it as Enterprise, Midrange or Lowrange.

Start the Qualification in the Task Menu with *Edit -> Controllers...* and select the Controller that has to be qualified.

Qualifying Disk Arrays

In Disk Arrays the Performance attributes of the Hardware (MDisks) that lies under it, are gathered. The gathered Data of the Disk Arrays can later be used at the generation of Reports (e.g. Tabel of all Raid 5 Arrays).

At the Storwize V7000, Disk Arrays are automatically created for each MDisk. No User-action is necessary. If one wants to summarize multiple MDisks into one Disk Array, one can discard the remaining Disk Array and create new ones.

You enter Disk Array Name, Raid Level, Disk type, number of disks (without Hotspares) and the Controller.
The Estimated I/O density (I/O Dichte) is calculated out of the specifications of the Raid Level and Disk type (max est. I/O Rate). This Value can in the case of deviation be overwritten.

Utilize the Treemap Path stored in Favorites for more efficient visualization of controllers, Disk Array and MDisk:
Favorites-> Predefined favorites -> Step by Step:Qualify and edit objects -> Define disk array for controllers.

Switch to the Treemap Path subsequently: BVQ -> Cluster -> Controller -> MDisk -> Disk array and switch to Object Color: BVQ qualified. MDisks that are not yet allocated to any Array, are then accompanied with the color Red.
You can directly in the Treemap create a new Disk Array for the Controllerand subsequently allocate the MDisks to the new Array by editting the Controller.

Qualifying Hosts

The Hosts are gathered via the BVQ scanning-procedure. For the Qualification the specification of the Room is required. Furthermore can the Comment Field be applied for Contact (above mentioned).

You start the Qualification in the Task Menu with *Edit -> Hosts* and select the Host that has to be qualified.

Qualifizierung SVC Nodes

The Nodes are also gathered via the BVQ scanning-procedure. For the Qualification the specification of the Room is required.

You start the Qualification in the Task Menu with *Edit -> Nodes* and select Node that has to be qualified.

Qualifying BVQ Accounting Package

The Accounting Package is part of the Premium Edition and is optionally available for the Performance Edition. You carry out the qualification-steps that are described in this collumn, if have the Premium Edition or the Performance Edition with the Accounting Package.

For the Accounting Package company-information is needed, and thereby the following questions are asked: "Which Applications are implemented in the company-environment?" or "How can these Applications be sensibly gathered in an "invoice-technical" way?
In BVQ it is for that matter necessary to summarise and expediently link the correspondent information subsequently.

- What are my cost units?
  E.g. Cost Unit: Human Resources Department, Customer xy or Software Development
  -> This Information is generally registered into the Group "Costcenter".

- Which Application Groups function under single Cost Units?
  E.g. Salary Accounting, centralized Chronology
  -> This Information is generally regitered into the Group "Application Group".

- Which Applications functions in the Application groups ?
  E.g. SAP Prod, SAP Test, Notes, Citrix or the development-software
  -> this Information is generally registered into the Group "Application".

- Are the Applications divided onto different Systems?
  E.g. Tablespaces, Backup
  -> This Information is generally registered into the Group "VDisk Group".

The VDisk Groups in this way display the Basis of the Accounting Package. With the VDisk Groups the SLA's, Data-Classes, Mobility Zones, VDisks and the VDisk Group Relations are connected.

Qualifying Cost Center, Application Group, Application, VDisk Group, VDisk

The Qualification is for instance meaningfully arranged, if one starts with the Cost Center and subsequently with the Application Groups, Application and finally carry out the qualification of the VDisk Groups. Thereby can the new object be connected with an already existing object.

Objekt Cost Centers

You start the Qualification in the Task Menu with *Edit -> Cost Centers...* and click on Add in order to create different Cost Centers.

Objekt Application groups

You start the Qualification in the Task Menu with *Edit -> Application Groups...* and click on Add in order to create the Application Groups. Allocate the suitable Cost Center to each Application Group.

Objekt Application

You start the Qualification in the Task Menu with *Edit -> Applications...* and click on Add in order to create the Applications. Allocate the suitable Application Group to each Application.

Objekt VDisk groups

You start the Qualification in the Task Menu with *Edit -> VDisk Groups...* and click on Add in order to create the VDisk Groups. Allocate the suitable Application to each VDisk Group.

VDisk

You start the Qualification in the Task Menu with *Edit -> VDisks...* and select a VDisk, which can be fittingly connected with a previously created VDisk Group.
Utilize the Name Filter in order to locate VDisks swiftly.

Then allocate the valid VDisk Group to it. Carry out the Qualification for the other VDisks as well.

After having these qualification steps, the first evaluations in the BVQ can take place (e.g. the capacity-view / Performance View on a Cost Center).

Capacity View:

Performance View:

Qualifying Mobility Zone

You now start to arrange the Mobility Zones. Mobility Zones determine, on which Storage Pools the VDisks of a VDisk Group should be placed. Thus, Regulations for files of VDisks are established, which then for instance can be inspected on their compliance (VDisk X should only be filed into the MDisk Groups Y & Z).

Create and allocating MDisk Groups
You start the Qualification in the Task Menu Bar with *Edit -> Mobility Zones...**and click on *Add**.

Allocating Mobility Zones to the VDisk Groups
You start the Qualification in Task Menu with *Edit -> VDisk groups...* and select a suitable passende VDisk Group.
A Mobility Zone can hereby be applied for multiple VDisk Groups.
Specify the Mobility Zone for the Mobilty Zone (P), in which the primary copies of the VDisks should be located in the VDisk Group.
Specify the Mobility Zone for the Mobilty Zone (S), in which the secondary copies of the VDisks should be located in the VDisk Group. Thus, mirrored VDisks (VDisk Mirror) can be allocated to different Mobility Zones.

1.12.2 Usage of the treemap for the optimization of storage costs

Objective:

Disks with moderate performance requirements should be moved from an expensive, high-performance storage pool to a favorable pool.

All VDisks with large capacity in the expensive storage pool are marked by choosing the aspect capacity. By pressing a button the aspect is changed to the aspect performance.

The markings are preserved and show all of the VDisks with their relative load values.

The VDisks with high load (large scaled) are deselected (with orange frame) and the remaining VDisks are moved by drag & drop into a favorable pool.

Abbildung 90 Optimizing memory costs: Screenshot high load VDisks

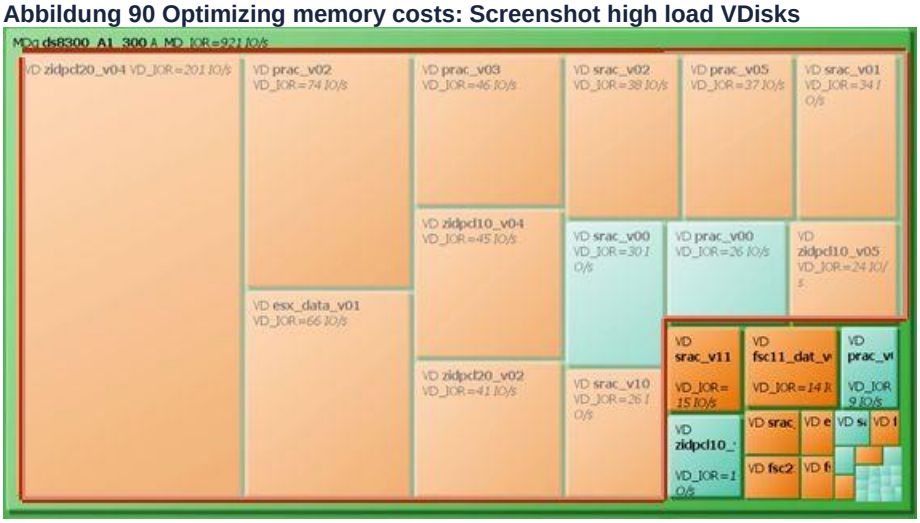
MDg ds8300_A1_300 A MD_IOR=971 IO/s
VD zidpcl20_v04 VD_IOR=201 IO/s
VD prac_v02 VD_IOR=74 IO/s
VD prac_v03 VD_IOR=46 IO/s
VD srac_v02 VD_IOR=38 IO/s
VD prac_v05 VD_IOR=37 IO/s
VD srac_v01 VD_IOR=34 I O/s
VD zidpcl10_v04 VD_IOR=45 IO/s
VD srac_v00 VD_IOR=30 I O/s
VD prac_v00 VD_IOR=26 IO/s
VD zidpcl10_v05 VD_IOR=24 IO/ s
VD esx_data_v01 VD_IOR=66 IO/s
VD srac_v11 VD_IOR= 15 IO/s
VD fsc11_dat_v VD_IOR=14 I
VD prac_v VD_IOR 9 IO/s
VD zidpcl20_v02 VD_IOR=41 IO/s
VD srac_v10 VD_IOR=26 I O/s
VD srac
VD zidpcl10_ VD_IOR=1 O/s
VD fsc2

The example highlights the identification of unfavourably placed datasets and their move to more favourably positions with just a few actions.

This way high performance and expensive disk space can quickly be freed from capacities, which can be provided with the same performance on more favorable memory.

Abbildung 91 Optimize storage costs - VDisks with high IO density

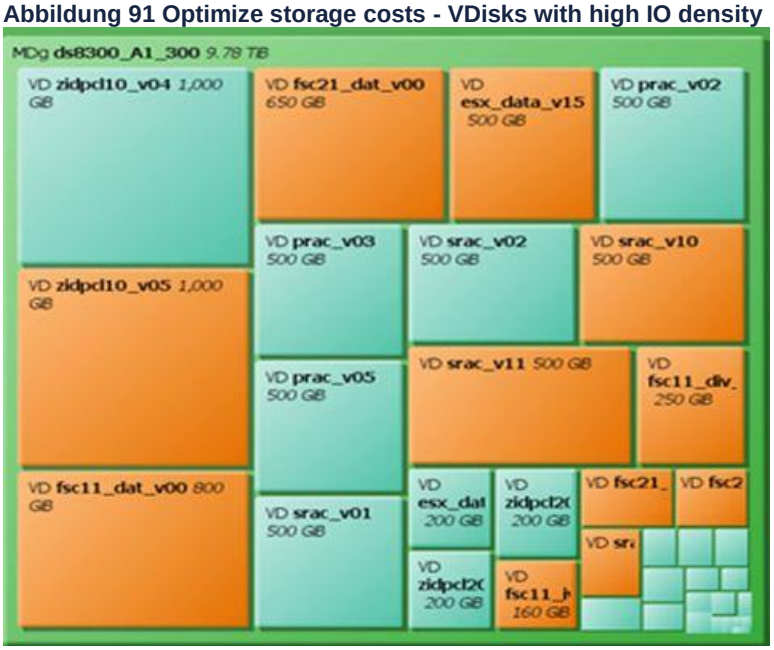

By deferring the viewing aspect to aspect capacity it gets obvious, how much space was occupied in the pool by non-performance volumes.

With the following drag and drop movement of marked volumes, the storage pool in this example was relieved by about 5TB.

1.12.3 Capacity efficiency

Evaluation of Thin Provisioning (space efficency) and compression.

Since SVC code version 6.4, Real-time Compression (RTC) is available in addition to the data reduction Thin Provisioning (space efficiency).

The purpose of this data reduction functionality is to significantly reduce the occupied space of a volume compared to its virtual capacity.

As a result, three types of volumes are now available:

1. Fully allocated volumes (FAV)
2. Thin provisioned / Space efficient vollumes (SEV)
3. Real-time compressed volumes (RCV) (based on SEV)

A key indicator for the savings-potential generated by the data reduction functions is the capacity efficiency. It expresses the size of the saved capacity.
This indicator is already crucial for thin provisioned VDisk copies. Their additional compression typically leads to a further increase of saved capacity .

The following BVQ functions were extended for evaluating the capacity efficiency:

- The chapter *accountables* of the object detail view
- Additional topology - color schemes

1.12.3.1 Extension of the object detail view

The object detail view contains two new paragraphs in the chapter accountables:

1. Compression
2. Total capacity efficiency

Abbildung 92 Object detail view with evaluation of the compression

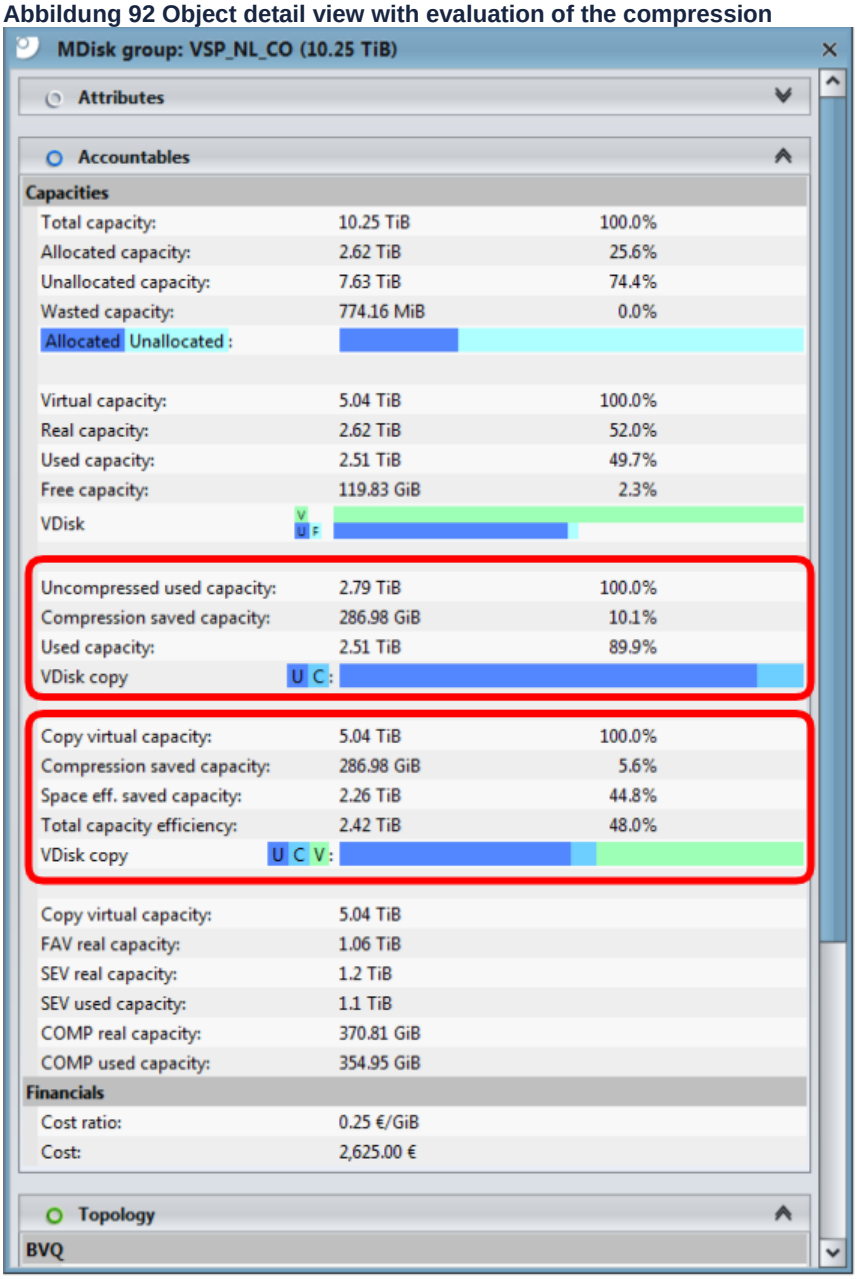
MDisk group: VSP_NL_CO (10.25 TiB)
Attributes
Accountables
Capacities
Total capacity: 10.25 TiB 100.0%
Allocated capacity: 2.62 TiB 25.6%
Unallocated capacity: 7.63 TiB 74.4%
Wasted capacity: 774.16 MiB 0.0%
Allocated Unallocated :
Virtual capacity: 5.04 TiB 100.0%
Real capacity: 2.62 TiB 52.0%
Used capacity: 2.51 TiB 49.7%
Free capacity: 119.83 GiB 2.3%
VDisk
Uncompressed used capacity: 2.79 TiB 100.0%
Compression saved capacity: 286.98 GiB 10.1%
Used capacity: 2.51 TiB 89.9%
VDisk copy U C :
Copy virtual capacity: 5.04 TiB 100.0%
Compression saved capacity: 286.98 GiB 5.6%
Space eff. saved capacity: 2.26 TiB 44.8%
Total capacity efficiency: 2.42 TiB 48.0%
VDisk copy U C V :
Copy virtual capacity: 5.04 TiB
FAV real capacity: 1.06 TiB
SEV real capacity: 1.2 TiB
SEV used capacity: 1.1 TiB
COMP real capacity: 370.81 GiB
COMP used capacity: 354.95 GiB
Financials
Cost ratio: 0.25 €/GiB
Cost: 2,625.00 €
Topology
BVQ

Paragraph compression

The paragraph *compression* contains the following metrics:

Metric	Calculation	Description
Uncompressed used capacity:	uncompressed_used_capacity	Theoretical size of the uncompressed used capacity

Compression saved capacity:	compression_saved_capacity = uncompressed_used_capacity - used_capacity	Compression gain: saved capacity by compression
Used capacity:	used_capacity	Compressed capacity
Bar chart: U Cs		Ratio of compressed (U) capacity to compression gain (Cs)

All percentage values are bases on the uncompressed size (Uncompressed used capacity = 100%).

Paragraph total capacity efficiency

Metric	Calculation		Description
Copy virtual capacity	copy_virtual_capacity	= SUM(vdiskcopy.virtual_capacity)	Sum of the virtual capacity of all VDisk copies
Space eff. saved capacity:	space_efficiency	= copy_virtual_capacity - uncompressed_used_capacity	Reduction benefit: saved capacity by Thin Provisioning
Compression saved capacity:	compression_efficiency	= compression_saved_capacity	Compression gain: saved capacity by compression
Total capacity efficiency	capacity_efficiency	= copy_virtual_capacity - allocated_capacity	Actually saved overall capacity
Bar chart: U Cs V			Ratio of compressed (U) capacity to compression gain (Cs) based on the virtual capacity

All percentage values are based on the sum of the virtual capacity of all VDisk copies (copy virtual capacity = 100%).

Note the difference between this simply calculated virtual capacity and the total virtual capacity:

- The simply calculated virtual capacity is the size of the volume which is visible from a server´s point of view.
- The total virtual capacity is calculated from the sum of the virtual capacity of all volume´s copies and represents the comparison value for the identification of saved capacity.

A simple example: a volume with a virtual capacity of 100GiB has 2 copies.

- Without Thin Provisioning or compression 200GiB memory would have to be allocated.
- If the compression is active for both copies and the compression rate has a vlue of 50%, 2 * 50 = 100GiB can be saved compared to a full occupancy (200GiB).

1.12.3.2 New color schemes for the treemap

Different topology - color schemes can be applied to the representation in the treemap.

Similarly to the metrics in the new paragraph total capacity efficiency new color schemes were added to visualize objects regarding their capacity efficiency.

The following color schemes are available:

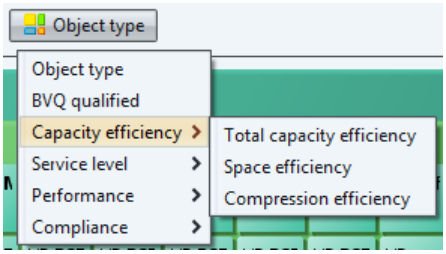

Color scheme	Based on metric
Total capacity efficiency	capacity_efficiency
Space efficiency	space_efficiency
Compression efficiency	compression_efficiency

Tabelle 78 Color scheme capacity efficiency - meaning of the colors

Color	Efficiency	Value range	Meaning
Red	Bad	0 - 24%	Poor compression rate or low data reduction by Thin Provisioning
Yellow	Average	25 - 49%	Average compression rate or mean data reduction by Thin Provisioning
Greeb	Good	50 - 79%	Good compression rate or higher data reduction by Thin Provisioning
Blue	Excellent	80 - 100%	Excellent compression rate or highest data reduction by Thin Provisioning
Faded	Not Applicable	--	Affected object type, but object does not possess the correct data-reduction type for the selected color

Object type	Not Applicable	--	Object type is not affected

Screenshots

Abbildung 93 Screenshot: staining

Abbildung 94 Screenshot: staining

Abbildung 95 Screenshot: staining

2 USER'S GUIDE - BVQ REPORTING

2.1 Reports overview

2.1.1 Standard reports

The standard reports are components of the BVQ performance- and premium edition.

The following standard reports are provided in BVQ:

- **Accounting report: STD-01_ACC_COST-CAP_ORGUN_DL1_FMT** - (Can only be applied with the accounting package) Overview about accounting units/applications and their utilisation of storage.
- **Historical Cluster Capacity Report: STD-02_CAP_REALCAP-LICCAP_VIRTU_DL1_FMT** - Historical evaluation of SVC / Storwize Systems.
- **Disksystem Capacity Report: STD-03_CAP-CAP-CNT_STGPR_DL1_FMT** - Evaluation of the connected disksystems.

2.1.2 Premium reports

The premium reports are components of the BVQ premium edition.

The following reports are additionally provided in BVQ premium Edition:

- **Accounting Report: Accounting_Costcenter-Application-VDiskGroup Report -** Overview about accounting units/ application-groups / applications / v-disk groups and their utilisation of storage.
- **Capacity Report: Capacitiy_StorageClass Report -** Evaluation of the connected disksystems and their storage performance classes.

2.1.3 Customer-orientated reports

Optionally, customer-orientated reports can be created for all editions. In cases of further queries, please contact your BVQ consultant.

⚠ BVQ entry edition does not contain reports!

2.2 Creating reports

By using the menu item **File-> start reporting** the selection window is opened.

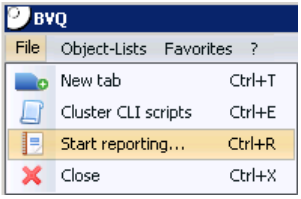

Afterwards click on the "browse" button and enter the path of the reports.

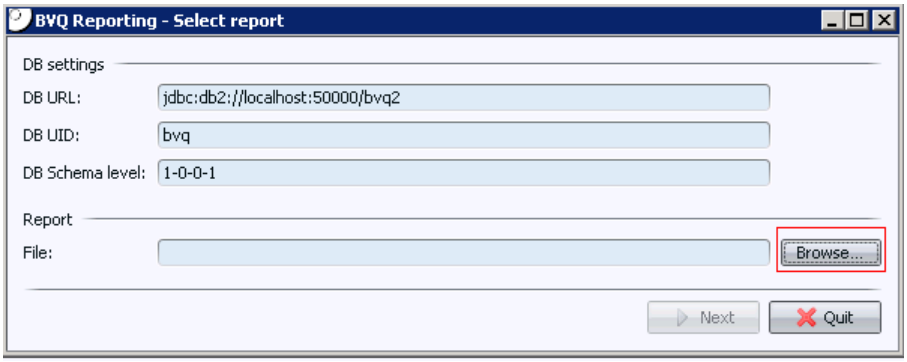

Select the appropriate report and click the "open" button.

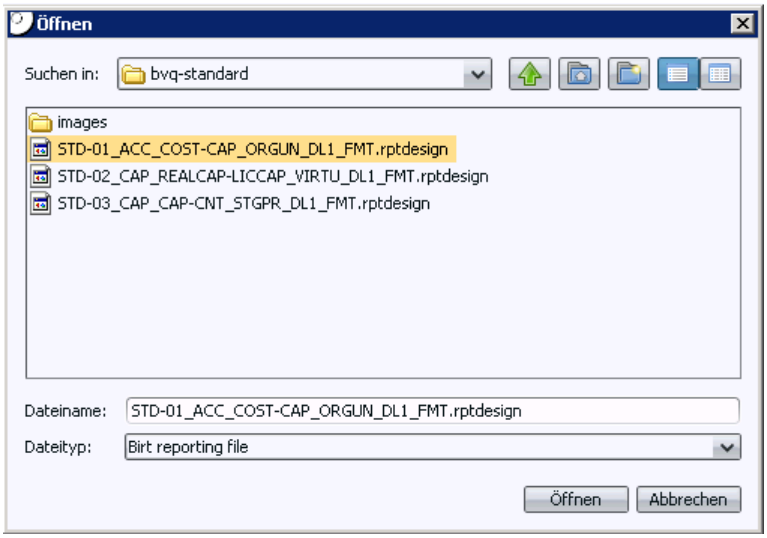

Click on "Save as..." and enter the path for saving. Different formats can be selected as file types. PDF is the preferred report format. Afterwards the report is created automatically. In large IT-environments this process can take a certain period of time.

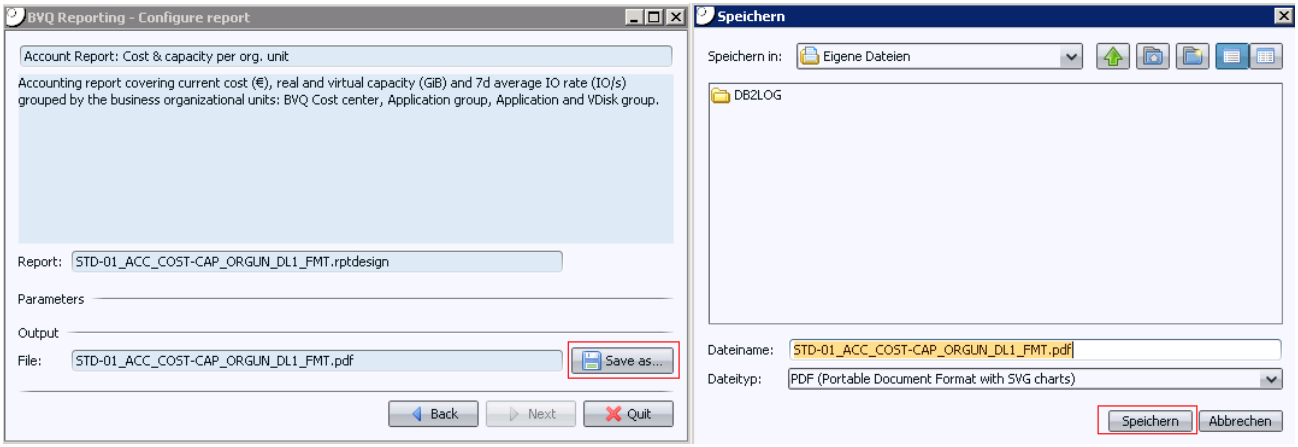
BVQ Reporting - Configure report
Account Report: Cost & capacity per org. unit
Accounting report covering current cost (€), real and virtual capacity (GiB) and 7d average IO rate (IO/s) grouped by the business organizational units: BVQ Cost center, Application group, Application and VDisk group.
Report: STD-01_ACC_COST-CAP_ORGUN_DL1_FMT.rptdesign
Parameters
Output
File: STD-01_ACC_COST-CAP_ORGUN_DL1_FMT.pdf
Save as...
Back
Next
Quit
Speichern
Speichern in: Eigene Dateien
DB2LOG
Dateiname: STD-01_ACC_COST-CAP_ORGUN_DL1_FMT.pdf
Dateityp: PDF (Portable Document Format with SVG charts)
Speichern
Abbrechen

After generating the report, it can be opened and evaluated by using a PDF reader.

Optionally custom reports are also possible, in which special adaptations can be implemented on the company's requirements.

2.3 ODBC access to the BVQ reporting scheme

2.3.1 Content

- Content
- Overview
- Download
- Installation
- Configuration
  - 1. Register DB2 ODBC driver in Windows
  - 2. Register data source in the Windows ODBC datasource administrator
- Procedure on 64 bit systems
- Usage in MS Excel 2007
- More information

2.3.2 Overview

For example to be able to access DB2 via MS Excel or OpenOffice an ODBC driver for DB2 must be installed.

On the BVQ DB server it is automatically installed and registered in the Microsoft ODBC. Download, installation and configuration steps 2.1 and 2.2 are not applicable in this case .
⚠ With DB2 9.7 FP4 the ODBC driver does not work completely with Windows 2008 R2 and Windows 7 64 bit. Please upgrade to DB2 9.7 FP6 (configuration problems with Odbcad).

Please also refer to the documentation of the BVQ report DB scheme.The DB views are documented there, which can be used to access the DB.

This is named at IBM: "IBM DB2 Driver for ODBC and CLI"

> CLI stands for call level interface. A C/C++ API that has nothing to do with the DB2 command line interface (`db2cmd`)
>
> Article about this in the DB2 Infocenter

2.3.3 Download

The "IBM DB2 driver for ODBC" can be download without charge from the Internet:

IBM Internet link

- 32 bit:
  - SVA Wiki (V9.7FP6)
- 64 bit:
  - SVA Wiki (V9.7FP6)

2.3.4 Installation

Unfotunately no installation program does exist for the IBM DB2 Driver for ODBC and CL I. It has to be installed manually:

1. Copy the ZIP archive containing the driver to the target system.
2. Unzip the TIP archive to a directory of your choice on the target machine (z.B. `%APPDATA%\IBM\db2cli`).
3. Optionally: delete the ZIP archive then.

2.3.5 Configuration

2.3.5.1 1. Register DB2 ODBC driver in Windows

```
D:\Programme\installable\IBM\DB2_Express\9.5\odbcdriver\bin>db2oreg1.
exe -h

 db2oreg1 - Install/Uninstall the IBM Data Server Driver for ODBC

 Usage: db2oreg1 -[i/u/s]

D:\Programme\installable\IBM\DB2_Express\9.5\odbcdriver\bin>db2oreg1.
exe -i

D:\Programme\installable\IBM\DB2_Express\9.5\odbcdriver\bin>
```

Unfortunately db2oreg1 doesn´t give any appropriate feedback i n case of a successful installation.

2.3.5.2 2. Register data source in the Windows ODBC datasource administrator

1. **Calling ODBCAD:**
   (Please also see procedure on 64 bit systems)
   a. On 64 bit systems for 32 bit ODBC clients (like e.g. Excel 2010 or Open Office 3.4)
      `C:\Windows\SysWOW64\odbcad32`
   b. On 64 bit systems for 64 bit ODBC clients
      `C:\Windows\system32\odbcad32`
   c. On 32 bit systems
      `odbcad32`
2. **Verifying DB2 driver registration – tab "driver"**

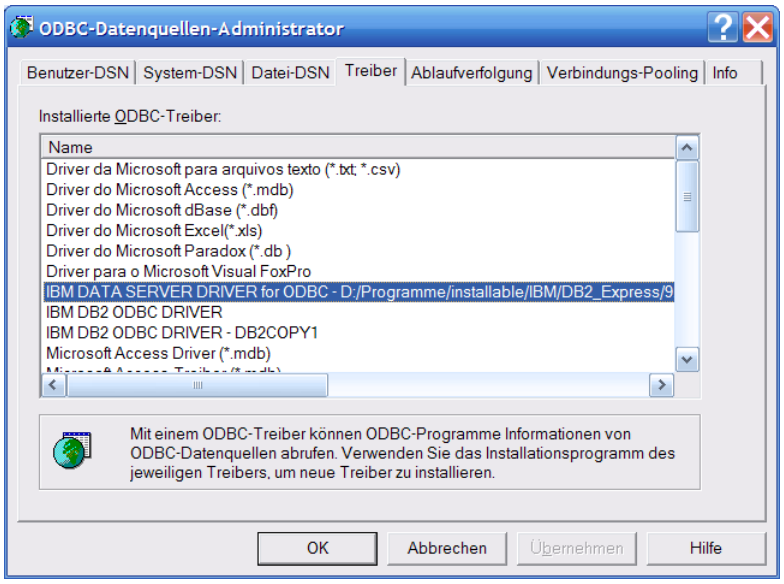

Show user data source - tab "user DSN" → "add"

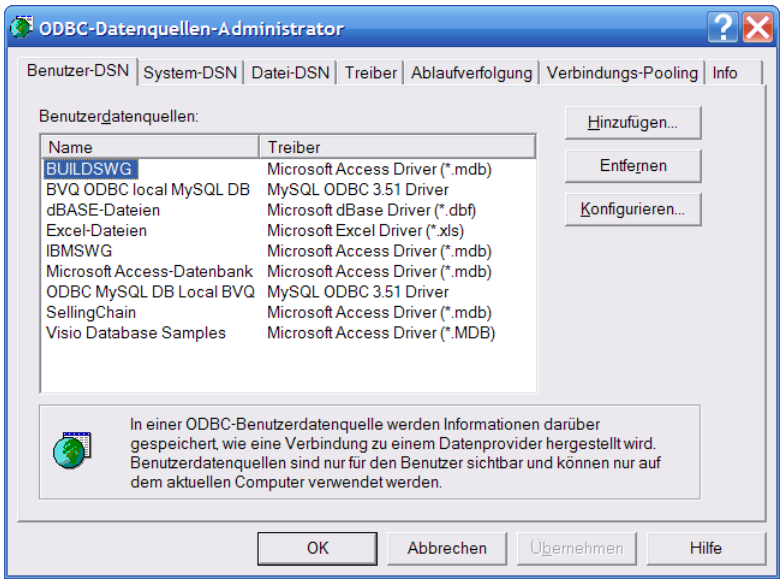

3. **Create new data source → "IBM DATA SERVER DRIVER for ODBC" → finish**

Neue Datenquelle erstellen
Wählen Sie einen Treiber aus, für den Sie eine Datenquelle erstellen möchten.
Name
Driver da Microsoft para arquivos texto (*.txt; *.csv)
Driver do Microsoft Access (*.mdb)
Driver do Microsoft dBase (*.dbf)
Driver do Microsoft Excel(*.xls)
Driver do Microsoft Paradox (*.db)
Driver para o Microsoft Visual FoxPro
IBM DATA SERVER DRIVER for ODBC - D:/Programme/ir
IBM DB2 ODBC DRIVER
IBM DB2 ODBC DRIVER - DB2COPY1
< Zurück
Fertig stellen
Abbrechen

4. **ODBC IBM DB2 Driver - add → data source name "BVQ"; database alias: "" → OK**

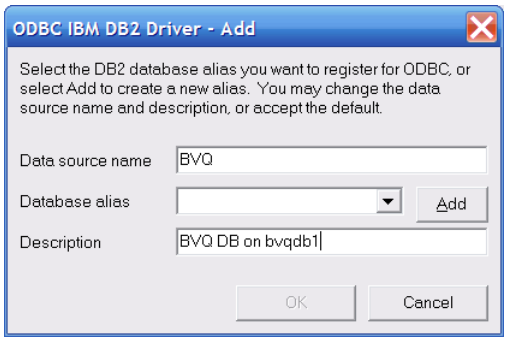
ODBC IBM DB2 Driver - Add
Select the DB2 database alias you want to register for ODBC, or select Add to create a new alias. You may change the data source name and description, or accept the default.
Data source name
BVQ
Database alias
Add
Description
BVQ DB on bvqdb1
OK
Cancel

5. **db2cli configure access parameters**
   a. Either via GUI:
   CLI/ODBC settings - BVQ → enter user ID and password → save password=√ → OK

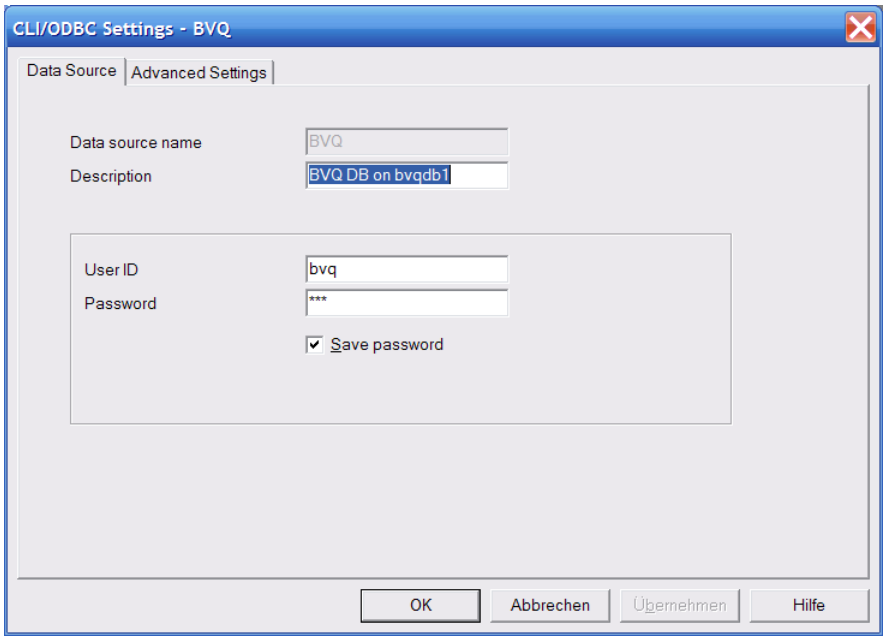

With "advanced settings db2cli.ini" all other parameters can be individually configured. The most important attribute is "hostname".

b. **Oder alternativ in der db2cli.ini:**
The db2cli.ini is either in the root directory of the unpacked ZIP file or in the directory documented in the IBM info center.
An example for a db2cli.ini:

```
; Comment lines start with a semi-colon.

[BVQ]
authentication=SERVER
database=bvq
protocol=TCPIP
port=50000
hostname=bvqdb1
uid=bvq
pwd=bvq
autocommit=0
TableType="'TABLE','VIEW','SYSTEM TABLE'"
SchemaList="'BVQ','BVQ_REPORT'"
```

6. **View user data source → BVQ → configure...**

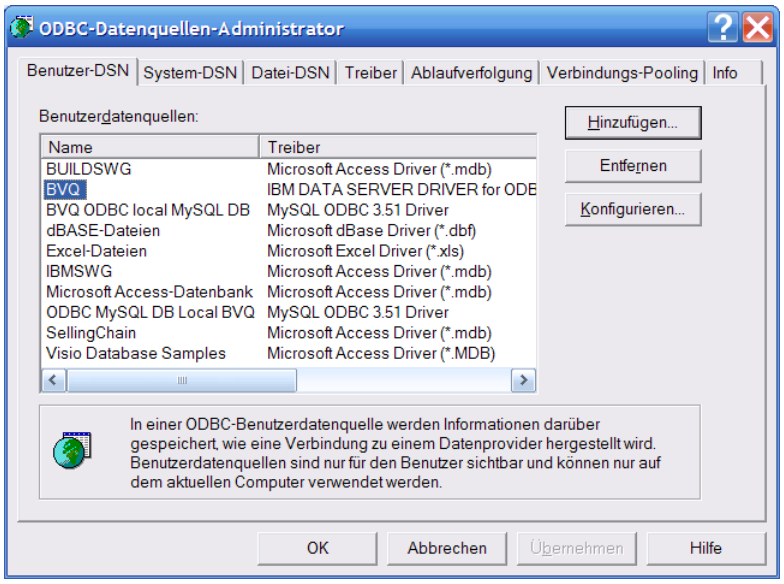

7. **Connection test: CLI/ODBC settings - BVQ → connect**

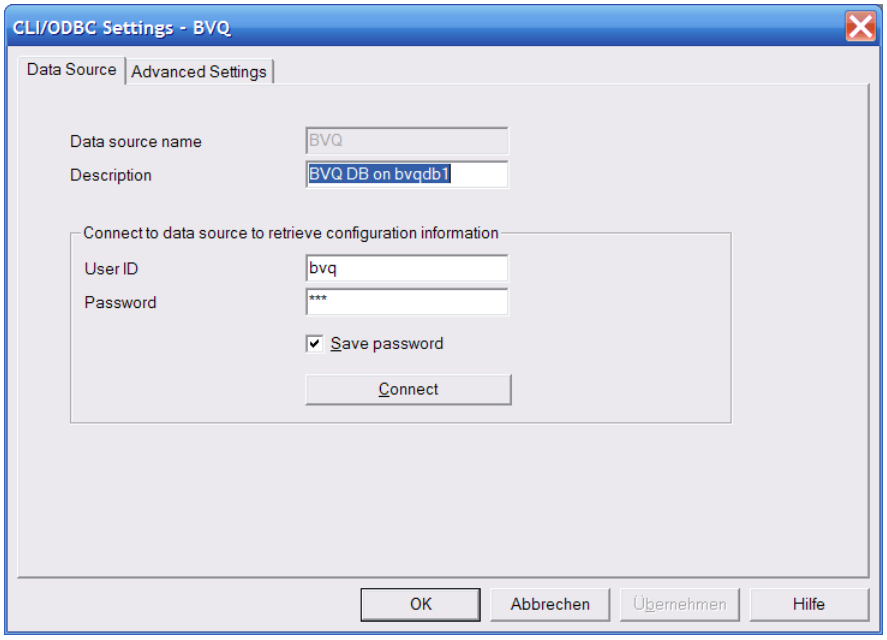

2.3.6 Procedure on 64 bit systems

On 64 bit systems a 32 bit ODBC is installed (Excel, OpenOffice etc. are 32 bit applications). odbcad32 must be called with C:\WINDOWS\SysWOW64\odbcad32.exe (then the 32 bit and not the 64 bit version of odbcad32 is truly called).

64 bit ODBC has only to be installed if a 64 bit application is used.

2.3.7 Usage in MS Excel 2007

1. **Excel -> data -> from other resources -> from Microsoft query**

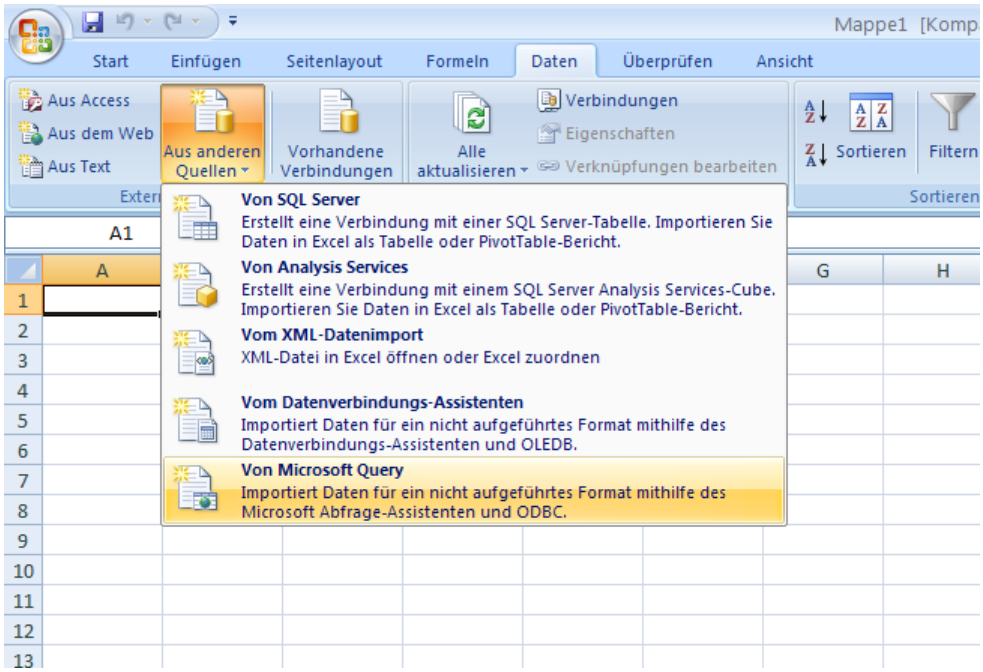

2. **Choose a data source -> BVQ***

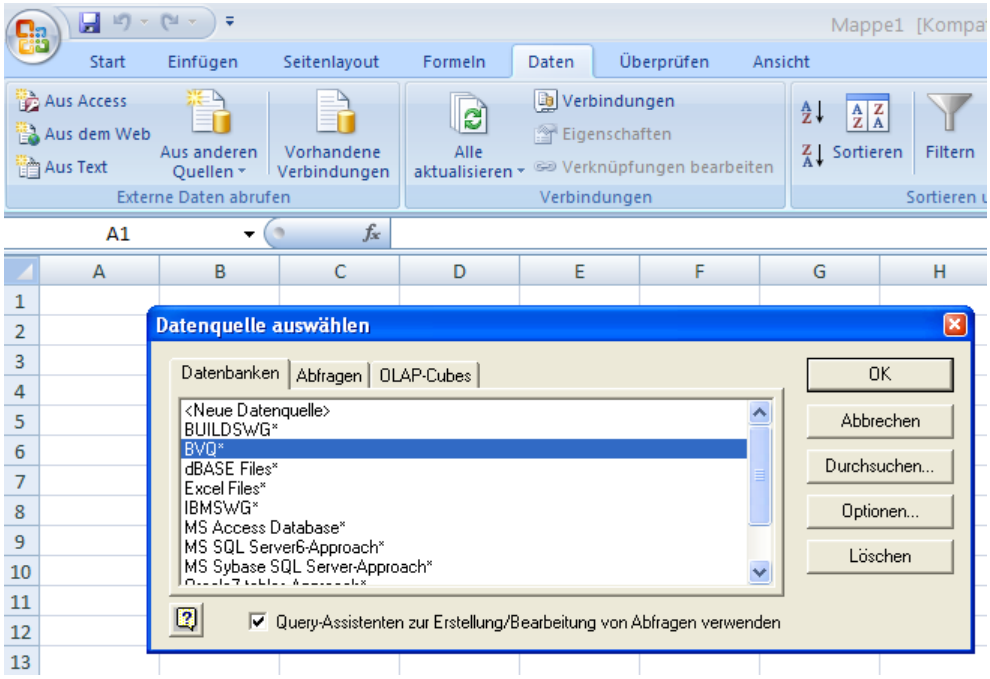

3. **Query wizard - choose columns**

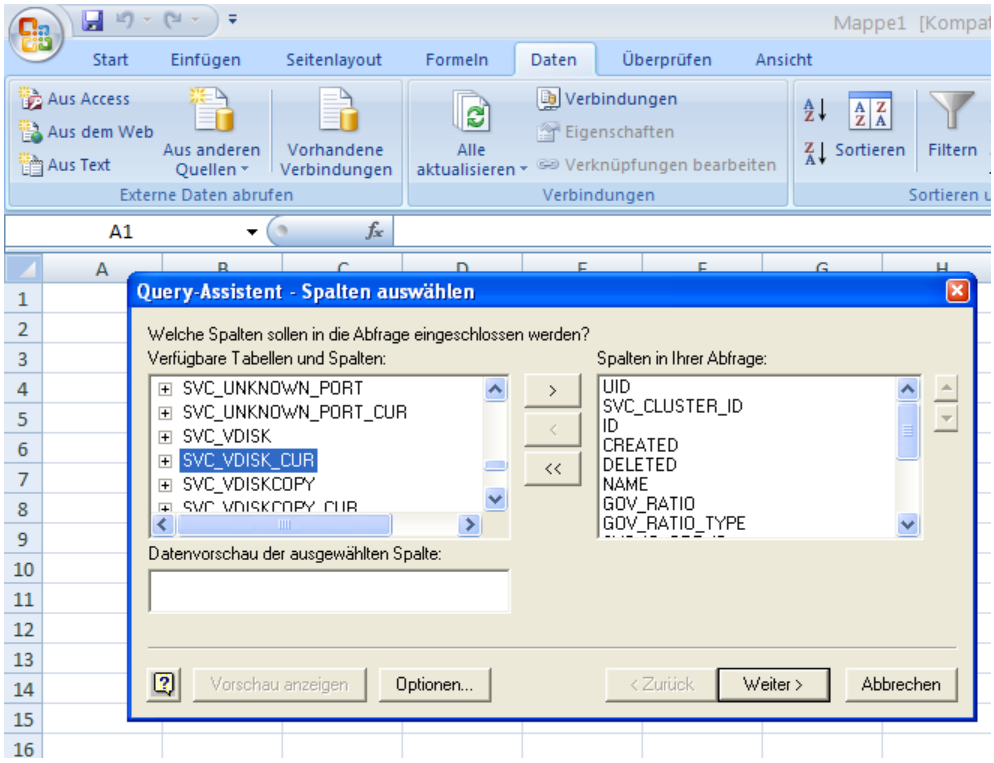

4. **Query wizard - filter data**

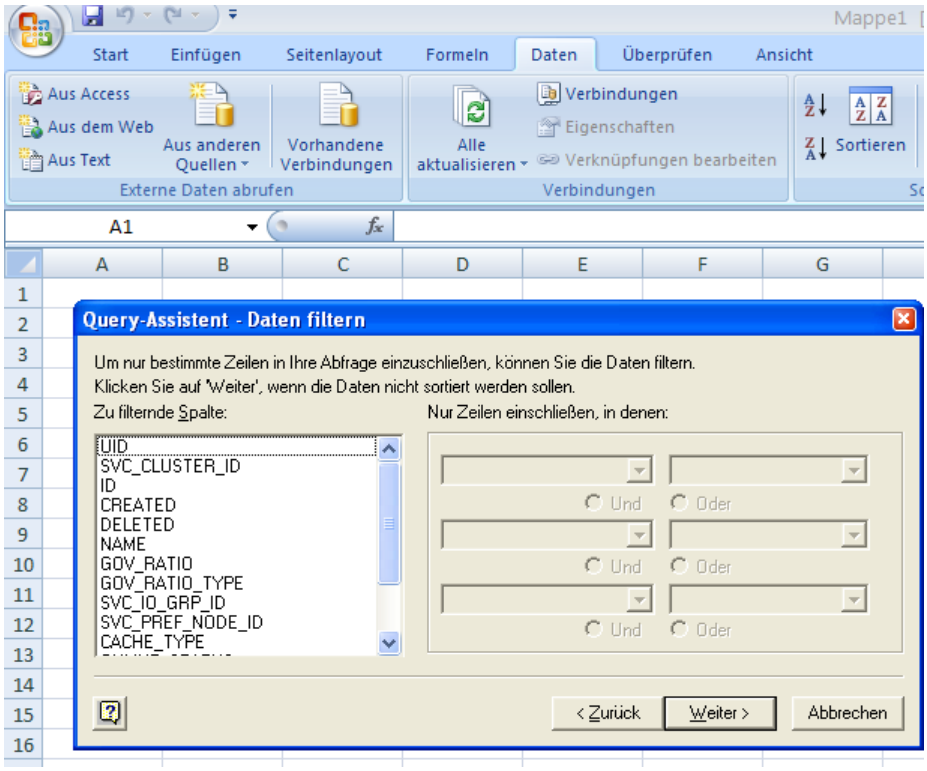

5. **Query wizard - finish -> return data to Microsoft Excel**

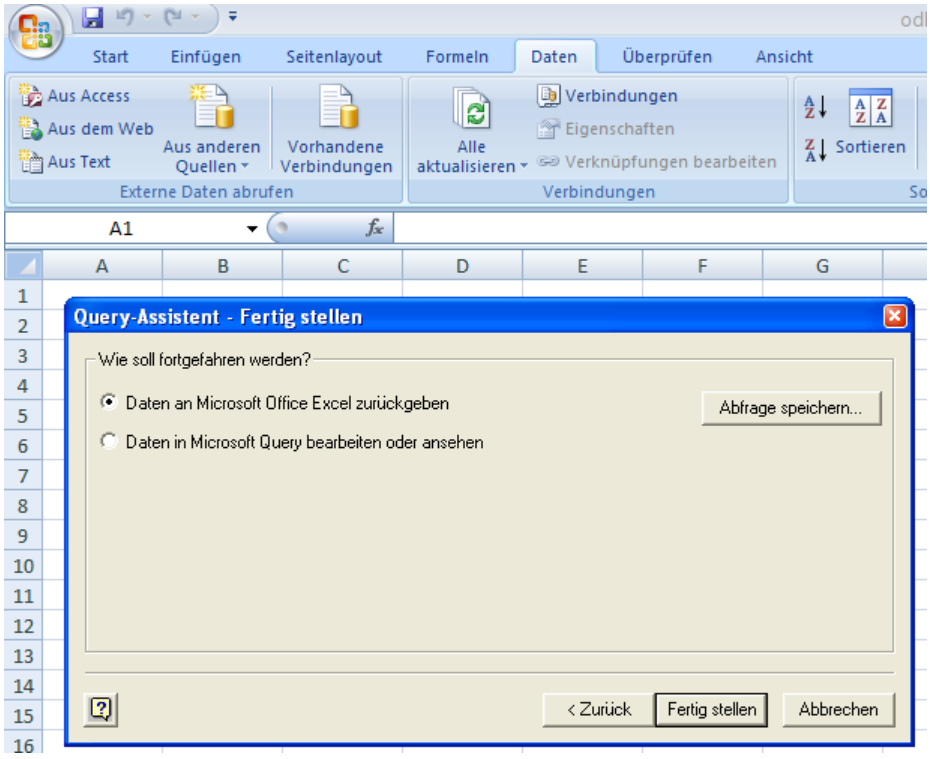

6. **Import data**

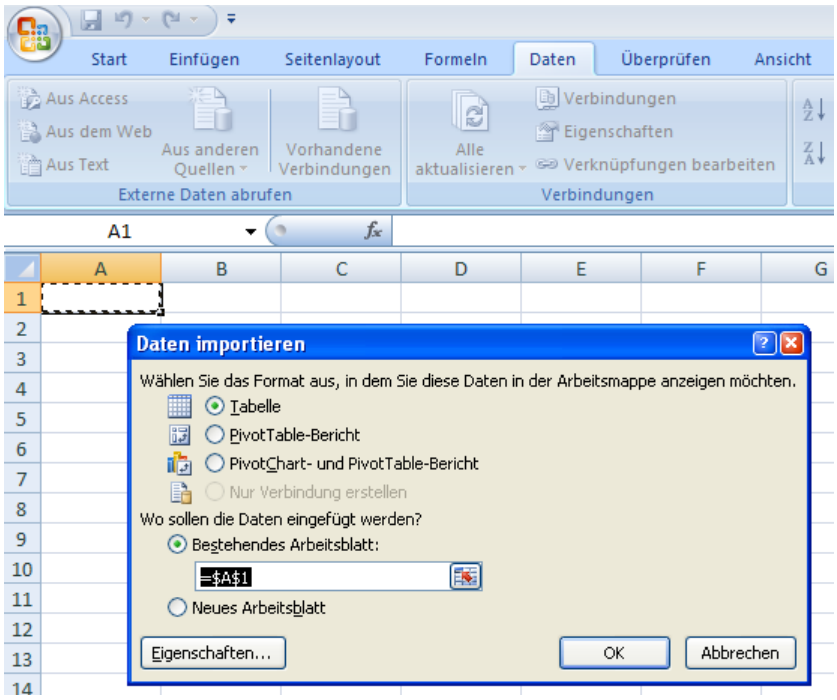

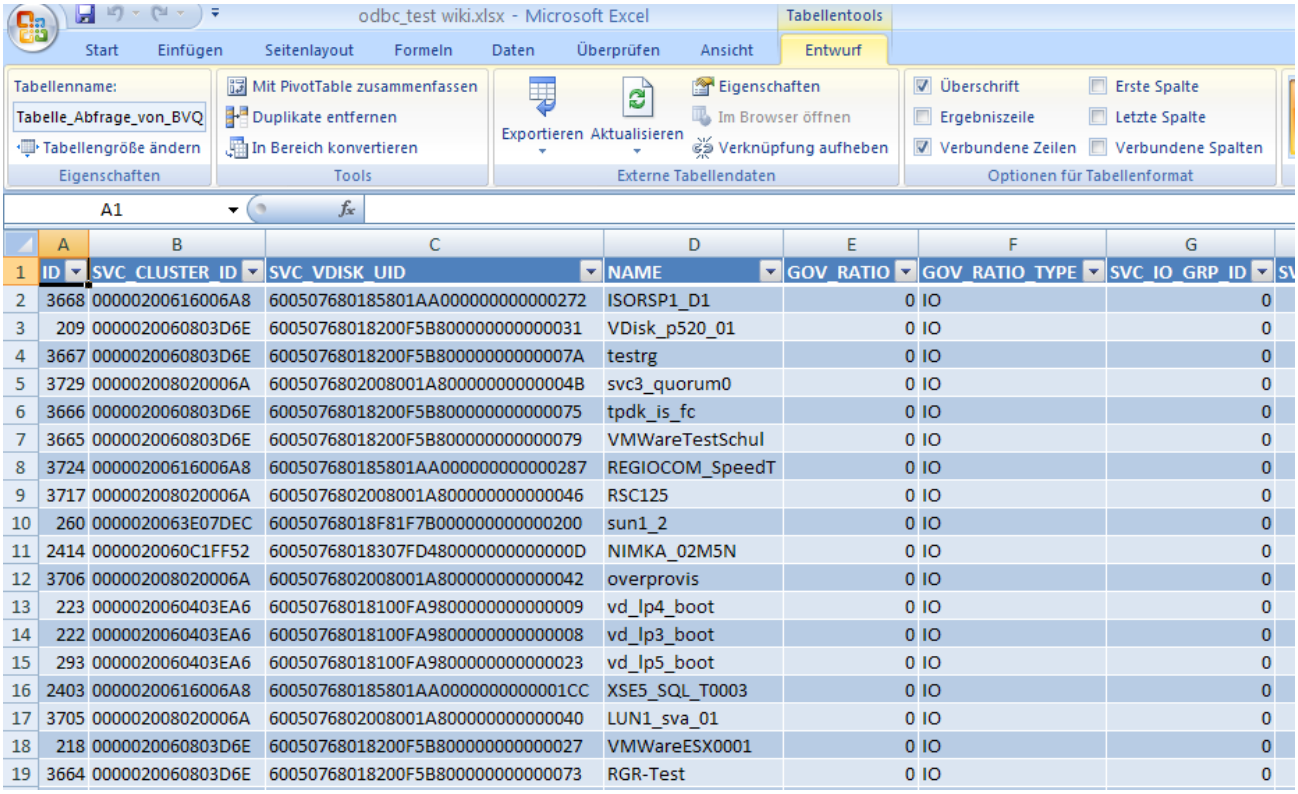

ID	SVC_CLUSTER_ID	SVC_VDISK_UID	NAME	GOV_RATIO	GOV_RATIO_TYPE	SVC_IO_GRP_ID
3668	0000020616006A8	600507680185801AA000000000000272	ISORSP1_D1	0	IO	0
209	0000020060803D6E	60050768018200F5B800000000000031	VDisk_p520_01	0	IO	0
3667	0000020060803D6E	60050768018200F5B80000000000007A	testrg	0	IO	0
3729	000002008020006A	6005076802008001A80000000000004B	svc3_quorum0	0	IO	0
3666	0000020060803D6E	60050768018200F5B800000000000075	tpdk_is_fc	0	IO	0
3665	0000020060803D6E	60050768018200F5B800000000000079	VMWareTestSchul	0	IO	0
3724	0000020616006A8	600507680185801AA000000000000287	REGIOCOM_SpeedT	0	IO	0
3717	000002008020006A	6005076802008001A80000000000046	RSC125	0	IO	0
260	0000020063E07DEC	60050768018F81F7B000000000000200	sun1_2	0	IO	0
2414	0000020060C1FF52	60050768018307FD480000000000000D	NIMKA_02M5N	0	IO	0
3706	000002008020006A	6005076802008001A80000000000042	overprovis	0	IO	0
223	0000020060403EA6	60050768018100FA9800000000000009	vd_lp4_boot	0	IO	0
222	0000020060403EA6	60050768018100FA9800000000000008	vd_lp3_boot	0	IO	0
293	0000020060403EA6	60050768018100FA9800000000000023	vd_lp5_boot	0	IO	0
2403	0000020616006A8	600507680185801AA0000000000001CC	XSE5_SQL_T0003	0	IO	0
3705	000002008020006A	6005076802008001A80000000000040	LUN1_sva_01	0	IO	0
218	0000020060803D6E	60050768018200F5B800000000000027	VMWareESX0001	0	IO	0
3664	0000020060803D6E	60050768018200F5B800000000000073	RGR-Test	0	IO	0

2.3.8 More information

- IBM DB2 Driver for ODBC and CLI overview

- Report DB scheme

2.4 BVQ database view scheme for report - scheme level 1.0.0.0.

BVQ-6926

The information has been compiled with the following statements:

Load all views in BVQ_REPORT

select name from sysibm.systables where CREATOR = 'BVQ_REPORT' AND type = 'V';
Load of the view-structure (here VDISK):
Select name,ColType,Length from Sysibm.syscolumns where tbname = 'VDISK_PERFORMANCE' AND TBCREATOR = 'BVQ_REPORT';
Load the view-structure (here VDISK) already in confluence wiki style:
Select CONCAT(CONCAT(CONCAT(CONCAT(CONCAT(CONCAT('|',name),'|'),ColType),'|'),Length),'|') from Sysibm.syscolumns where tbname = 'VDISK' AND TBCREATOR = 'BVQ_REPORT';

APP_GROUP_HIST_1_YEAR

COST_CENTER_ID	"INTEGER "	4
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
ID	"INTEGER "	4
MARKER_DATE	TIMESTMP	10
NAME	"VARCHAR "	64
NOTES	"VARCHAR "	2048

APP_HIST_1_YEAR

APP_GRP_ID	"INTEGER "	4
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
ID	"INTEGER "	4
MARKER_DATE	TIMESTMP	10
NAME	"VARCHAR "	64

NOTES	"VARCHAR "	2048

CLUSTER

SVC_CLUSTER_ID	"CHAR "	16
SVC_CLUSTER_NAME	"VARCHAR "	64

CLUSTER_CAPACITY

SVC_ALLOCATABLE_CAPACITY	"BIGINT "	8
SVC_ALLOCATED_CAPACITY	"BIGINT "	8
SVC_CLUSTER_ID	"CHAR "	16
SVC_CLUSTER_NAME	"VARCHAR "	64
SVC_LICENSE_FLASH_ENTITLED	"DOUBLE "	8
SVC_LICENSE_FLASH_USED	"DOUBLE "	8
SVC_LICENSE_REMOTE_ENTITLED	"DOUBLE "	8
SVC_LICENSE_REMOTE_USED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_ENTITLED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_USED	"DOUBLE "	8
SVC_UNALLOCATED_CAPACITY	"BIGINT "	8

CLUSTER_CAPACITY_HISTORY

SVC_ALLOCATABLE_CAPACITY	"BIGINT "	8
SVC_ALLOCATED_CAPACITY	"BIGINT "	8
SVC_CLUSTER_CHANGED	TIMESTMP	10
SVC_CLUSTER_DELETED	SMALLINT	2
SVC_CLUSTER_ID	"CHAR "	16
SVC_CLUSTER_NAME	"VARCHAR "	64
SVC_LICENSE_FLASH_ENTITLED	"DOUBLE "	8
SVC_LICENSE_FLASH_USED	"DOUBLE "	8
SVC_LICENSE_REMOTE_ENTITLED	"DOUBLE "	8
SVC_LICENSE_REMOTE_USED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_ENTITLED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_USED	"DOUBLE "	8
SVC_UNALLOCATED_CAPACITY	"BIGINT "	8

CLUSTER_CAPACITY_HISTORY_1_YEAR

MARKER_DATE	TIMESTMP	10
SVC_ALLOCATABLE_CAPACITY	"BIGINT "	8
SVC_ALLOCATED_CAPACITY	"BIGINT "	8
SVC_CLUSTER_CHANGED	TIMESTMP	10
SVC_CLUSTER_DELETED	SMALLINT	2
SVC_CLUSTER_ID	"CHAR "	16
SVC_CLUSTER_NAME	"VARCHAR "	64
SVC_LICENSE_FLASH_ENTITLED	"DOUBLE "	8
SVC_LICENSE_FLASH_USED	"DOUBLE "	8

SVC_LICENSE_REMOTE_ENTITLED	"DOUBLE "	8
SVC_LICENSE_REMOTE_USED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_ENTITLED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_USED	"DOUBLE "	8
SVC_UNALLOCATED_CAPACITY	"BIGINT "	8

CLUSTER_CAPACITY_HISTORY_2_YEARS

MARKER_DATE	TIMESTMP	10
SVC_ALLOCATABLE_CAPACITY	"BIGINT "	8
SVC_ALLOCATED_CAPACITY	"BIGINT "	8
SVC_CLUSTER_CHANGED	TIMESTMP	10
SVC_CLUSTER_DELETED	SMALLINT	2
SVC_CLUSTER_ID	"CHAR "	16
SVC_CLUSTER_NAME	"VARCHAR "	64
SVC_LICENSE_FLASH_ENTITLED	"DOUBLE "	8
SVC_LICENSE_FLASH_USED	"DOUBLE "	8
SVC_LICENSE_REMOTE_ENTITLED	"DOUBLE "	8
SVC_LICENSE_REMOTE_USED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_ENTITLED	"DOUBLE "	8
SVC_LICENSE_VIRTUALIZATION_USED	"DOUBLE "	8
SVC_UNALLOCATED_CAPACITY	"BIGINT "	8

COST_CENTER_HIST_1_YEAR

CREATED	TIMESTMP	10
DELETED	SMALLINT	2
ID	"INTEGER "	4
MARKER_DATE	TIMESTMP	10
NAME	"VARCHAR "	64
NOTES	"VARCHAR "	2048

DISK_ARRAY

BVQ_ARRAY_ID	"INTEGER "	4
BVQ_ARRAY_NAME	"VARCHAR "	64
BVQ_DISK_COUNT	"INTEGER "	4
BVQ_DISK_TYPE_NAME	"VARCHAR "	56
BVQ_IO_DENSITY	"DOUBLE "	8
BVQ_RAID_MODE	"VARCHAR "	4

DISK_TYPE

BVQ_CAPACITY	"BIGINT "	8
BVQ_CLASS	"VARCHAR "	10
BVQ_DISK_TYPE_ID	"INTEGER "	4
BVQ_DISK_TYPE_NAME	"VARCHAR "	56
BVQ_IO_RATE	"INTEGER "	4

BVQ_RPM	SMALLINT	2

HOST_CAPACITY

BVQ_HOST_ID	"INTEGER "	4
SVC_ALLOCATED_CAPACITY	"BIGINT "	8
SVC_CLUSTER_ID	"CHAR "	16
SVC_CLUSTER_NAME	"VARCHAR "	64
SVC_FREE_CAPACITY	"BIGINT "	8
SVC_HOST_ID	"INTEGER "	4
SVC_HOST_NAME	"VARCHAR "	64
SVC_REAL_CAPACITY	"BIGINT "	8
SVC_USED_CAPACITY	"BIGINT "	8
SVC_VIRTUAL_CAPACITY	"BIGINT "	8

HOST_PORT

BVQ_HOST_ID	INTEGER	4
SVC_CLUSTER_ID	CHAR	16
SVC_CLUSTER_NAME	VARCHAR	64
SVC_HOST_ID	INTEGER	4
SVC_WWPN	VARCHAR	128

MDISK

BVQ_DISK_ARRAY_ID	INTEGER	4
BVQ_MDISK_ID	INTEGER	4
SVC_CLUSTER_ID	CHAR	16
SVC_CONTROLLER_ID	INTEGER	4
SVC_MDISK_GRP_ID	INTEGER	4
SVC_MDISK_ID	INTEGER	4
SVC_MDISK_NAME	VARCHAR	64
SVC_MODE	VARCHAR	10

MDISK_CAPACITY

BVQ_ARRAY_ID	INTEGER	4
BVQ_MDISK_ID	INTEGER	4
SVC_ALLOCATABLE_CAPACITY	BIGINT	8
SVC_ALLOCATED_CAPACITY	BIGINT	8
SVC_CLUSTER_ID	CHAR	16
SVC_CONTROLLER_ID	INTEGER	4
SVC_EXTENT_COUNT	INTEGER	4
SVC_EXTENT_SIZE	INTEGER	4
SVC_FREE_EXTENTS	INTEGER	4
SVC_MDISK_ID	INTEGER	4
SVC_MDISK_NAME	VARCHAR	64
SVC_MODE	VARCHAR	10

SVC_UNALLOCATED_CAPACITY	BIGINT	8
SVC_UNMANAGED_CAPACITY	BIGINT	8
SVC_UNMANAGED_COUNT	INTEGER	4

MDISK_GROUP_CAPACITY

ALLOCATABLE_CAPACITY	BIGINT	8
ALLOCATED_CAPACITY	BIGINT	8
SVC_CLUSTER_ID	CHAR	16
SVC_MDISK_GROUP_ID	INTEGER	4
SVC_MDISK_GROUP_NAME	VARCHAR	64
UNALLOCATED_CAPACITY	BIGINT	8

MDISK_GROUP_CAPACITY_HISTORY

ALLOCATABLE_CAPACITY	BIGINT	8
BVQ_MDISK_GROUP_ID	INTEGER	4
SVC_CLUSTER_ID	CHAR	16
SVC_MDISK_GROUP_CHANGED	TIMESTMP	10
SVC_MDISK_GROUP_DELETED	SMALLINT	2
SVC_MDISK_GROUP_ID	INTEGER	4
SVC_MDISK_GROUP_NAME	VARCHAR	64
UNALLOCATED_CAPACITY	BIGINT	8

SVC_CLUSTER_HIST_1_YEAR

ALLOCATED_CAPACITY	BIGINT	8
CODE_LEVEL	VARCHAR	63
CONSOLE_IP	VARCHAR	21
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
ID	CHAR	16
IP_ADDRESS	VARCHAR	15
LICENSE_COMPRESSION_CAPACITY	DOUBLE	8
LICENSE_FLASH	DOUBLE	8
LICENSE_REMOTE	DOUBLE	8
LICENSE_VIRTUALIZATION	DOUBLE	8
MARKER_DATE	TIMESTMP	10
NAME	VARCHAR	64
REAL_CLUSTER_ID	CHAR	16
SECOND_IP_ADDRESS	VARCHAR	15
SERVICE_IP_ADDRESS	VARCHAR	15
STATS_FREQUENCY	INTEGER	4
STATS_STATUS	VARCHAR	3
TIMEZONE	VARCHAR	40
TOTAL_FREE_SPACE	BIGINT	8

TOTAL_MDISKGRP_CAPACITY	BIGINT	8
TOTAL_MDISK_CAPACITY	BIGINT	8
USED_COMPRESSION_CAPACITY	DOUBLE	8
USED_FLASH	DOUBLE	8
USED_REMOTE	DOUBLE	8
USED_VIRTUALIZATION	DOUBLE	8

SVC_CLUSTER_HIST_2_YEARS

ALLOCATED_CAPACITY	BIGINT	8
CODE_LEVEL	VARCHAR	63
CONSOLE_IP	VARCHAR	21
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
ID	CHAR	16
IP_ADDRESS	VARCHAR	15
LICENSE_COMPRESSION_CAPACITY	DOUBLE	8
LICENSE_FLASH	DOUBLE	8
LICENSE_REMOTE	DOUBLE	8
LICENSE_VIRTUALIZATION	DOUBLE	8
MARKER_DATE	TIMESTMP	10
NAME	VARCHAR	64
REAL_CLUSTER_ID	CHAR	16
SECOND_IP_ADDRESS	VARCHAR	15
SERVICE_IP_ADDRESS	VARCHAR	15
STATS_FREQUENCY	INTEGER	4
STATS_STATUS	VARCHAR	3
TIMEZONE	VARCHAR	40
TOTAL_FREE_SPACE	BIGINT	8
TOTAL_MDISKGRP_CAPACITY	BIGINT	8
TOTAL_MDISK_CAPACITY	BIGINT	8
USED_COMPRESSION_CAPACITY	DOUBLE	8
USED_FLASH	DOUBLE	8
USED_REMOTE	DOUBLE	8
USED_VIRTUALIZATION	DOUBLE	8

SVC_EXTENTS_HIST_1_YEAR

COPY_ID	SMALLINT	2
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
EXTENT_COUNT	INTEGER	4
MARKER_DATE	TIMESTMP	10
SVC_CLUSTER_ID	CHAR	16

SVC_MDISK_ID	INTEGER	4
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_UID	CHAR	32

SVC_MDISK_GROUP_HIST_1_YEAR

CAPACITY	BIGINT	8
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
EASY_TIER	VARCHAR	4
EASY_TIER_STATUS	VARCHAR	8
EXTENT_SIZE	INTEGER	4
ID	INTEGER	4
MARKER_DATE	TIMESTMP	10
NAME	VARCHAR	64
ONLINE_STATUS	VARCHAR	16
SVC_CLUSTER_ID	CHAR	16
UNALLOCATED_CAPACITY	BIGINT	8
WARNING	SMALLINT	2

SVC_VDISKCOPY_HIST_1_YEAR

AUTOEXPAND	SMALLINT	2
COPY_ID	SMALLINT	2
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
EASY_TIER	VARCHAR	3
EASY_TIER_STATUS	VARCHAR	8
FREE_CAPACITY	BIGINT	8
GRAINSIZE	SMALLINT	2
MARKER_DATE	TIMESTMP	10
ONLINE_STATUS	VARCHAR	10
REAL_CAPACITY	BIGINT	8
SVC_CLUSTER_ID	CHAR	16
SVC_MDISK_GRP_ID	INTEGER	4
SVC_VDISKCOPY_ROLE_ID	INTEGER	4
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_UID	CHAR	32
SYNC	SMALLINT	2
TYPE	VARCHAR	10
UNCOMPRESSED_USED_CAPACITY	BIGINT	8
USED_CAPACITY	BIGINT	8
WARNING	SMALLINT	2

SVC_VDISK_HIST_1_YEAR

CACHE_TYPE	VARCHAR	10
CAPACITY	BIGINT	8
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
FORMATTED	SMALLINT	2
GOV_RATIO	INTEGER	4
GOV_RATIO_TYPE	VARCHAR	2
ID	INTEGER	4
MARKER_DATE	TIMESTMP	10
NAME	VARCHAR	64
ONLINE_STATUS	VARCHAR	10
SVC_CLUSTER_ID	CHAR	16
SVC_IO_GRP_ID	INTEGER	4
SVC_PREF_NODE_ID	INTEGER	4
SYNC_RATE	SMALLINT	2
UID	CHAR	32

TIMEBAR_1_YEAR

MARKER_DATE	TIMESTMP	10

TIMEBAR_2_YEARS

MARKER_DATE	TIMESTMP	10

TOPOLOGY_CLUSTER__SITE__ROOM__NODE

BVQ_NODE_ID	INTEGER	4
BVQ_ROOM_ID	INTEGER	4
BVQ_ROOM_NAME	VARCHAR	64
BVQ_SITE_ID	INTEGER	4
BVQ_SITE_NAME	VARCHAR	64
SVC_CLUSTER_ID	CHAR	16
SVC_CLUSTER_NAME	VARCHAR	64
SVC_NODE_ID	INTEGER	4
SVC_NODE_NAME	VARCHAR	64

TOPOLOGY_COST_CENTER__APPLICATION_GROUP__APPLICATION__VDISK_GROUP

BVQ_APPLICATION_GROUP_ID	INTEGER	4
BVQ_APPLICATION_GROUP_NAME	VARCHAR	64
BVQ_APPLICATION_ID	INTEGER	4
BVQ_APPLICATION_NAME	VARCHAR	64
BVQ_COST_CENTER_ID	INTEGER	4
BVQ_COST_CENTER_NAME	VARCHAR	64
BVQ_VDISK_GROUP_ID	INTEGER	4
BVQ_VDISK_GROUP_NAME	VARCHAR	64

TOPOLOGY_SITE__ROOM__CONTROLLER__CLUSTER__ARRAY__MDISK

BVQ_ARRAY_ID	INTEGER	4
BVQ_ARRAY_NAME	VARCHAR	64
BVQ_CONTROLLER_ID	INTEGER	4
BVQ_CONTROLLER_MERGE_ID	INTEGER	4
BVQ_CONTROLLER_MERGE_NAME	VARCHAR	64
BVQ_MDISK_ID	INTEGER	4
BVQ_ROOM_ID	INTEGER	4
BVQ_ROOM_NAME	VARCHAR	64
BVQ_SITE_ID	INTEGER	4
BVQ_SITE_NAME	VARCHAR	64
SVC_CLUSTER_ID	CHAR	16
SVC_CLUSTER_NAME	VARCHAR	64
SVC_CONTROLLER_ID	INTEGER	4
SVC_CONTROLLER_NAME	VARCHAR	64
SVC_MDISK_ID	INTEGER	4
SVC_MDISK_NAME	VARCHAR	64

VDISK

BVQ_VDISK_GROUP_ID	INTEGER	4
BVQ_VDISK_ID	INTEGER	4
SVC_CLUSTER_ID	CHAR	16
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	VARCHAR	64
SVC_VDISK_UID	CHAR	32

VDISKCOPY_CAPACITY

BVQ_VDISKCOPY_ID	INTEGER	4
SVC_ALLOCATED_CAPACITY	BIGINT	8
SVC_CLUSTER_ID	CHAR	16
SVC_FREE_CAPACITY	BIGINT	8
SVC_MDISK_GRP_ID	INTEGER	4
SVC_REAL_CAPACITY	BIGINT	8
SVC_USED_CAPACITY	BIGINT	8
SVC_VDISKCOPY_ID	SMALLINT	2
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	VARCHAR	64
SVC_VDISK_UID	CHAR	32
SVC_VDISK_VIRTUAL_CAPACITY	BIGINT	8

VDISKCOPY_CAPACITY_HIST_1_YEAR

BVQ_VDISKCOPY_ID	INTEGER	4
MARKER_DATE	TIMESTMP	10
SVC_ALLOCATED_CAPACITY	BIGINT	8

SVC_CLUSTER_ID	CHAR	16
SVC_FREE_CAPACITY	BIGINT	8
SVC_MDISK_GRP_ID	INTEGER	4
SVC_REAL_CAPACITY	BIGINT	8
SVC_USED_CAPACITY	BIGINT	8
SVC_VDISKCOPY_ID	SMALLINT	2
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	VARCHAR	64
SVC_VDISK_UID	CHAR	32
SVC_VDISK_VIRTUAL_CAPACITY	BIGINT	8

VDISKCOPY_FINANCIAL

BVQ_VDISKCOPY_ID	INTEGER	4
SVC_CLUSTER_ID	CHAR	16
SVC_COST	DECIMAL	31
SVC_VDISKCOPY_ID	SMALLINT	2
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	VARCHAR	64
SVC_VDISK_UID	CHAR	32

VDISKCOPY_HIST_1_YEAR

AUTOEXPAND	SMALLINT	2
COPY_ID	SMALLINT	2
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
GRAINSIZE	SMALLINT	2
ID	INTEGER	4
MARKER_DATE	TIMESTMP	10
REAL_CAPACITY	BIGINT	8
SVC_CLUSTER_ID	CHAR	16
SVC_MDISK_GRP_ID	INTEGER	4
SVC_VDISKCOPY_ROLE_ID	INTEGER	4
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_UID	CHAR	32
WARNING	SMALLINT	2

VDISK_CAPACITY

BVQ_VDISK_ID	INTEGER	4
SVC_ALLOCATED_CAPACITY	BIGINT	8
SVC_CLUSTER_ID	CHAR	16
SVC_FREE_CAPACITY	BIGINT	8
SVC_REAL_CAPACITY	BIGINT	8
SVC_USED_CAPACITY	BIGINT	8

SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	VARCHAR	64
SVC_VDISK_UID	CHAR	32
SVC_VIRTUAL_CAPACITY	BIGINT	8

VDISK_FINANCIAL

SVC_CLUSTER_ID	CHAR	16
SVC_COST	DECIMAL	31
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	VARCHAR	64
SVC_VDISK_UID	CHAR	32

VDISK_GROUP_HIST_1_YEAR

APPLICATION_ID	INTEGER	4
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
ID	INTEGER	4
MARKER_DATE	TIMESTMP	10
NAME	VARCHAR	64
NOTES	VARCHAR	2048

VDISK_HIST_1_YEAR

CAPACITY	BIGINT	8
CREATED	TIMESTMP	10
DELETED	SMALLINT	2
GOV_RATIO	INTEGER	4
GOV_RATIO_TYPE	VARCHAR	2
ID	INTEGER	4
MARKER_DATE	TIMESTMP	10
NAME	VARCHAR	64
SVC_CLUSTER_ID	CHAR	16
SVC_IO_GRP_ID	INTEGER	4
SVC_PREF_NODE_ID	INTEGER	4
SVC_VDISK_UID	CHAR	32
VDISK_GRP_ID	INTEGER	4

VDISK_PERFORMANCE

BVQ_ID	INTEGER	4
SVC_CLUSTER_ID	CHAR	16
SVC_VDISK_ID	INTEGER	4
SVC_VDISK_NAME	CHAR	32
SVC_VDISK_UID	CHAR	32
T3_IND01	DOUBLE	8

3 ANNEX

3.1 Most important links

Interesting links for further information

Web link directory:

- BVQ Wiki (http://bvqwiki.sva.de)
- BVQ Homepage (http://www.bvq-software.com)